arXiv:1612.06404v1 [stat.ME] 19 Dec 2016

RANDOM WALK MODELS OF NETWORK FORMATION AND SEQUENTIAL MONTE CARLO METHODS FOR GRAPHS

Benjamin Bloem-Reddy and Peter Orbanz

Columbia University

We introduce a class of network models that insert edges by connecting the starting and terminal vertices of a random walk on the network graph. Within the taxonomy of statistical network models, this class is distinguished by permitting the location of a new edge to explicitly depend on the structure of the graph, but being nonetheless statistically and computationally tractable. In the limit of infinite walk length, the model converges to an extension of the preferential attachment model—in this sense, it can be motivated alternatively by asking what preferential attachment is an approximation to. Theoretical properties, including the limiting degree sequence, are studied analytically. If the entire history of the graph is observed, parameters can be estimated by maximum likelihood. If only the final graph is available, its history can be imputed using MCMC. We develop a class of sequential Monte Carlo algorithms that are more generally applicable to sequential random graph models, and may be of interest in their own right. The model parameters can be recovered from a single graph generated by the model. Applications to data clarify the role of the random walk length as a length scale of interactions within the graph.

1. Introduction. We consider problems in network analysis where the observed data constitutes a single graph. Tools developed for such problems include, for example, graphon models [8, 34, 25, 7, 2, 47, 19], the subfamily of stochastic blockmodels [e.g. 22], preferential attachment graphs [4, 18, 43], interacting particle models [32], and others. In this context, a statistical model is a family $\mathcal{P} = \{P_\theta, \theta \in \mathbf{T}\}$ of probability distributions on graphs, and sample data is explained either as a single graph drawn from the model, or as a subgraph of such a graph.

The class of models studied in the following generate a graph by inserting one edge at a time: Start with a single edge (with its two terminal vertices). For each new edge, select a vertex V in the current graph at random.

- With a fixed probability $\alpha \in (0, 1]$, add a new vertex and connect it to V.
- Otherwise, start a random walk at V, and connect its terminal vertex to V.

In contrast to the models listed at the outset, the location of newly inserted edges depends on those of previously inserted ones, through the random walk. The dependence results in some unusual properties. One is that, with two scalar parameters, the model generates a perhaps unexpected range of different graph structures (see Figure 1 in Section 2 for examples). The random walk step can be regarded as a model of interactions:

- Long walks can connect vertices far apart in the graph. In this sense, the (expected) length of the walk is a range scale.

- Where convenient, the walk may be interpreted explicitly, as a process on the network (e.g. users in a social network forming connections through other users). More generally, it biases connections towards vertices reachable along multiple paths.

Thus, random walk models can explain a range of graph structures as the outcome of interactions on a certain length scale. As the ensuing discussion will explain in detail, the models are statistically and analytically tractable.

Background. At a conceptual level, the work presented here is motivated by the question how statistical network models handle stochastic dependence within a network; the discussion in this paragraph is purely heuristic. Consider the link graph of a network, regarded as a system of random variables (representing edges and vertices). Informally, for structure to occur in a network with non-negligible probability, the variables must be dependent. The graphs below are (i) a protein-protein interaction network (see Section 5), and (ii) its reconstruction sampled from a graphon model [8, 34] fitted to (i):

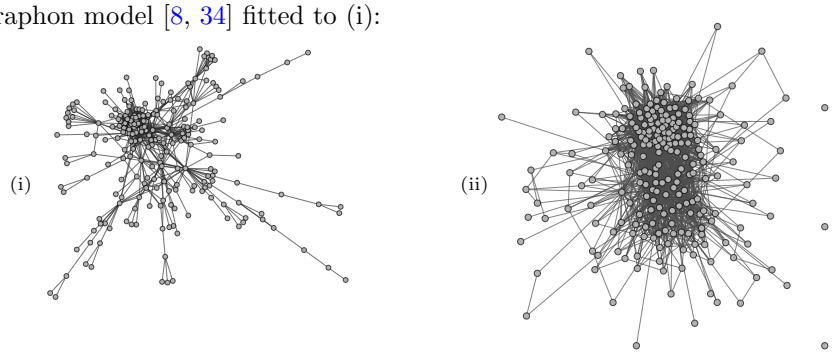

The data set (i) contains pendants (several degree-1 vertices arranged around a single vertex), long isolated chains of edges, etc. For these to arise in a random graph requires local dependence between edges on different length-scales, and they are conspicuously absent in the graph on the right. That is not a shortcoming of the method used to fit the model, but inherent to graphons per se, since these models constrain edges to be conditionally independent given certain vertex-wise information. If the missing structure is relevant to statistical analysis, its absence constitutes a form of model misspecification. To put the problem into context, it is useful to compare some commonly used network models:

- Graphon models (as mentioned above). These have recently received considerable attention in statistics [e.g. 25, 7, 2, 47, 19], and include stochastic blockmodels [22], and many models popular in the machine learning literature [see 38].
- Models that randomize a graph to match a given degree sequence, such as the configuration model [18].
- Models describing a formation process, such as preferential attachment (PA) models [4], the CHKNS model [18], or interacting particle models [32].

All of these models constrain dependence by rendering edges conditionally inde-

pendent given some latent variable: In a graphon model, this latent variable is a local edge density; in the related model proposed in [12], a scalar variable associated with each vertex; in a configuration model, the degree sequence; in preferential attachment models, the degree sequence history of the current graph at each step of the formation process. The models above have all been noted to be misspecified for many network analysis problems:

- In the case of graphon models, various symptoms of this problem have been noted [e.g. 9, 38]: The generated graph is dense; unless it is k-partite or disconnected, the distance between any two vertices in an infinite graph is almost surely 1 or 2; etc. A graphon can encode any fixed pattern on some number n of vertices, but this pattern then occurs on every possible subgraph of size n with fixed probability.
- Configuration models are popular in probability (due to their simplicity), but have limited use in statistics unless the quantity of interest is the degree sequence itself: They are "maximally random" given the degrees, in a manner similar to exponential families being maximally random given a sufficient statistic, and are thus insensitive to any structure not captured by the degrees.
- Preferential attachment models have been thoroughly studied in applied probability, since they explain a commonly observed, non-trivial property of networks (power law degrees) by a simple mechanism (the preferential attachment rule) [43]. As statistical models, they lack expressiveness: The basic PA model has only a single, integer-valued parameter, the number of edges generated by each new vertex.

In each instance of misspecification listed, the constraints on dependence are culpable. Nonetheless, such constraints are imposed for good reasons. One is tractability. As in any system of multiple random variables, forms of stochastic dependence amenable to mathematical analysis and statistical inference are rare. Another is the basic design principle that a statistical model should be as simple as possible, and that it is arguably the presence, not the absence, of dependence that requires justification.

The approach taken here is to make a modeling assumption on the network formation process (as e.g. preferential attachment models do); whether that assumption is considered adequate or not must depend on the problem at hand. Conceptually, however, the random walk models studied in the following allow slightly more intricate forms of stochastic dependence than the models discussed above, without sacrificing applicability.

Article overview. Random walk models are defined in Section 2. They can be chosen to generate either multigraphs or simple graphs, which in either case are sparse. A quantity that plays a key role in both theory and inference is the history of graph, i.e. the order Π in which edges are inserted. In dynamic networks [e.g. 29, 18], where a graph is observed over time, Π is an observed variable. If only the final graph is observed, Π is latent. Section 3 establishes some theoretical properties:

- If the entire history Π of a multigraph is observed, model parameters can be

estimated by maximum likelihood (Section 3.2).
- Under suitable conditions, the limiting degree distribution can be characterized (Theorem 3.2). The generated graphs can exhibit power law properties.
- Conditionally on the order in which vertices are inserted, the normalized degree sequence converges to an almost sure limit. Results of this type are known for preferential attachment models (for which the vertex order is deterministic, and hence need not be conditioned upon). Each limiting relative degree can be generated marginally by a sampling scheme reminiscent of "stick-breaking". See Theorem 3.3.
- If the length of the random walk tends to infinity, the model converges to a generalization of the preferential attachment model (Proposition 3.4).

If only the final graph is observed, inference is still possible, by treating the history Π as a latent variable and imputing it using a sampling algorithm. This problem is of broader interest, since it permits the application of sequential or dynamic network models, including the PA model and many others, to a single observed graph. In Section 4, we show how the graph history can be sampled, under any sequential network model satisfying a Markov and a monotonicity property, using a Sequential Monte Carlo (SMC) algorithm for random graphs. Using this algorithm, we derive inference algorithms for sequential network models that can be designed as SMC or particle MCMC methods [e.g. 3, 17]. In Section 5, we discuss empirical results and the application of the model to data:

- Tractability of the model is be evaluated by performing parameter inference on graphs generated by the model. The model parameters can indeed be recovered from a single observed graph; see Section 5.1.
- The role of the random walk parameter can be understood in more detail by relating it to the empirical mixing time of simple random walk on the data set; see Section 5.2.
- Comparisons between the random walk model and other network models are given in Section 5.3; we also discuss some issues raised by such comparisons.

The latent order variable Π can be related to certain measures of vertex centrality, which is briefly illustrated on data in Section 5.4.

2. Random walk graph models: Definition. A **random walk** on a connected graph g started at vertex v is a random sequence $(v, V_1, \ldots, V_k)$ of vertices, such that neighbors in the sequence are connected in g. Figuratively, a walker repeatedly moves along an edge to a randomly selected neighbor, until k edges have been crossed. Edges and vertices may be visited multiple times. The random walk is **simple** if the next vertex is selected with uniform probability from the neighbors of the current one. The distribution of the terminal vertex V_k of a simple random walk of length k on g is denoted $\mathrm{SRW}(g, v, k)$. The model considered in the following is, like many sequential network models, most easily specified in terms of a generative algorithm:

Algorithm 1 (Random walk model on multigraphs)**.**

- Fix $\alpha \in (0, 1]$, a distribution P on $\mathbb{N}$, and a graph G_1 with a single edge connecting two vertices.
- For $t = 2, \ldots, T$, generate G_t from G_{t-1} as follows:
  (1) Select a vertex V of G_{t-1} at random (see below).
  (2) With probability α, attach a new vertex to V.
  (3) Else, draw $K \sim P$ and connect V to $V' \sim \mathrm{SRW}(G_{t-1}, V, K)$.

The vertex V in (1) is either sampled uniformly from the current vertex set, or from the **size-biased** (or **degree-biased**) distribution $\mathbb{S}(v) = \deg(v)/\sum_{v' \in \mathcal{V}(g)} \deg(v')$, where $\deg(v)$ is the degree of vertex v in g.

Algorithm 1 generates multigraphs, since the two vertices selected by the random walk may already be connected (resulting in a multi-edge), or may coincide (resulting in a self-loop). To generate simple graphs instead, step (3) above is replaced by:

(3') Else, draw $K \sim P$ and $V' \sim \mathrm{SRW}(G_{t-1}, V, K)$. Connect V to V' if they are distinct and not connected; else, attach a new vertex to V.

Either sampling scheme defines the law of a sequence $(G_1, G_2, \ldots)$ of graphs. We denote this law $\mathbf{RW}_{\mathrm{U}}(\alpha, P)$ if the vertex V in (1) is chosen uniformly, or $\mathbf{RW}_{\mathrm{SB}}(\alpha, P)$ in the size-biased case. Each such law is defined both on multigraphs and on simple graphs, depending on which sampling scheme is used.

We generally choose the length of the random walk as a Poisson variable (although most theoretical results in the next section hold for any choice of P): Denote by $\mathrm{Poisson}_+(\lambda)$ the Poisson distribution with parameter λ, shifted by 1, i.e. the law of $K + 1$, where K is Poisson. We write $\mathbf{RW}(\alpha, \lambda)$ for $\mathbf{RW}(\alpha, \mathrm{Poisson}_+(\lambda))$ where convenient. Examples of graphs generated by this distribution are shown in Figure 1.

3. Model properties. We first observe that graphs generated by the model are sparse by construction. Let $(G_1, \ldots, G_t, \ldots)$ be a graph sequence generated by any **RW** model on simple or multigraphs, with parameter α, where G_1 has a single edge connecting two vertices. Denote by $\mathcal{V}(G_t)$ and $\mathcal{E}(G_t)$ the set of vertices and of edges, respectively, in G_t. For any **RW** model, each vertex is added with one edge and there is a constant positive probability at each step of adding a new vertex, which leads to the following:

Observation. *Graphs generated by the random walk model are sparse: In a sequence* $(G_1, G_2, \ldots) \sim \mathbf{RW}(\alpha, P)$*, the number of edges grows as* $|\mathcal{E}(G_t)| = \Theta(|\mathcal{V}(G_t)|)$.

3.1. *Mixed random walks and the graph Laplacian.* Most results in the following involve the law of a random walk. For a walk of fixed length, this law is determined

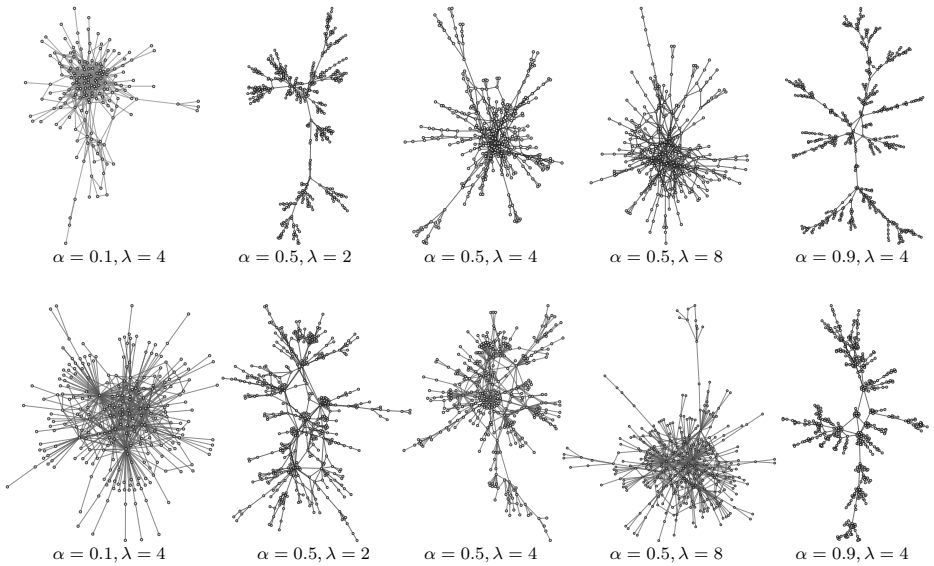
$\alpha = 0.1, \lambda = 4$
$\alpha = 0.5, \lambda = 2$
$\alpha = 0.5, \lambda = 4$
$\alpha = 0.5, \lambda = 8$
$\alpha = 0.9, \lambda = 4$
$\alpha = 0.1, \lambda = 4$
$\alpha = 0.5, \lambda = 2$
$\alpha = 0.5, \lambda = 4$
$\alpha = 0.5, \lambda = 8$
$\alpha = 0.9, \lambda = 4$

FIG 1. *Examples of simple graphs generated by a random walk model:* $\mathbf{RW}_{\mathrm{U}}(\alpha, \mathit{Poisson}_+(\lambda))$ *distribution (top row), and by a* $\mathbf{RW}_{\mathrm{SB}}(\alpha, \mathit{Poisson}_+(\lambda))$ *distribution (bottom row).*

by the graph's Laplacian matrix [15]. If the length is randomized, one has to mix against its distribution; for suitable choice of the length distribution, the mixed law of the random walk is still available in closed form.

Consider an undirected graph G with n vertices, adjacency matrix A and degree matrix $D = \mathrm{diag}(\deg(v_1), \ldots, \deg(v_n))$. The probability that a simple random walk started at a vertex u terminates at v after exactly k steps is the entry (u, v) of the matrix $(D^{-1}A)^k$. If the length of the walk is a random variable K with law P, the marginal probability of reaching v from u is hence

$$\mathbb{P}\{V_{end} = v | V_0 = u,\, G\} \quad = \quad \Big[\sum_{k\in\mathbb{N}} (D^{-1}A)^k P\{K = k\}\Big]_{uv} \; = \big[\mathbb{E}_P\left(D^{-1}A\right)\big]_{uv} \; . \tag{1}$$

A useful way to represent the information in $D^{-1}A$ is as the matrix

$$\mathbf{L} := D^{1/2}(\mathbb{I}_n - D^{-1}A)D^{-1/2} \qquad \text{with } \mathbb{I}_n := \text{identity matrix} \, , \tag{2}$$

known as the **normalized graph Laplacian** [15]. The law of a random walk of random length is obtained as follows:

PROPOSITION 3.1. *Let g be a connected, undirected graph on n vertices, and define*

$$\mathbf{K}^{\lambda} := e^{-\lambda \mathbf{L}} \qquad \textit{and} \qquad \mathbf{K}^{r,p} := \big(\mathbb{I}_n + \tfrac{p}{1-p}\mathbf{L}\big)^{-r} \, . \tag{3}$$

Then for a simple random walk $(V_0, \dots, V_K)$ *of random length* K *on* g*,*

$$\mathbb{P}\{V_{end} = v | V_0 = u,\, g\} = \left[D^{-1/2}(\mathbb{I}_n - \mathbf{L})\mathbf{K}D^{1/2}\right]_{uv}, \tag{4}$$

with $\mathbf{K} = \mathbf{K}^{\lambda}$ *if* K *has law* $\text{Poisson}_+(\lambda)$*, or* $\mathbf{K} = \mathbf{K}^{r,p}$ *if* K *has law* $\text{NB}_+(r,p)$*.*

For a Poisson length, the result derives from the relationship

$$\sum_{k=0}^{\infty} e^{-\lambda} \frac{\lambda^k t^k}{k!} B^k = e^{-\lambda(\mathbb{I}_n - tB)} \qquad \text{for any } B \in \mathbb{R}^{n\times n} \tag{5}$$

of the Poisson convolution semigroup to the matrix exponential function. It extends to the negative binomial distribution, since the latter is obtained from the Poisson by mixing: Let $\text{NB}_+(r,p)$ denote the negative binomial distribution, again shifted to positive integers, with parameter r and p. The random walk model satisfies

$$\mathbf{RW}_{\text{U}}(\alpha, \text{NB}_+(r,p)) = \int \mathbf{RW}_{\text{U}}(\alpha, \text{Poisson}_+(\lambda))\; \gamma(\lambda \,|\, r, \tfrac{p}{1-p}) d\lambda \tag{6}$$

where $\gamma(\,\bullet \,|\, a, b)$ denotes the gamma density with parameters a and b. The same holds for $\mathbf{RW}_{\text{SB}}$. The matrices (3) arise in other contexts, in particular in the machine learning literature: The **heat kernel** $\mathbf{K}^{\lambda}$ and the **regularized Laplacian kernel** $\mathbf{K}^{r,p}$ have applications in collaborative filtering, semi-supervised learning and manifold learning [30, 31]. Both act as a smoothing operators on functions defined on the vertex set, by damping the eigenvalue spectrum [41].

3.2. *Maximum likelihood estimation for fully observed sequences.* We now consider parameter estimation in the special case of multigraphs in the "dynamic" case, where the entire history of a graph is observed. Inference in general is covered in Section 4. Let $(G_1, \dots, G_t)$ be a multigraph sequence generated by any **RW** model, with parameter α, where G_1 has a single edge connecting two vertices. Since the model creates vertices by independent α-coin flips, the number N_t of vertices in G_t is $(N_t - 2) \sim \text{Binomial}(t-1, \alpha)$, and $\hat{\alpha}_t := (N_t - 2)/(t-1)$ is a maximum likelihood estimator for α. Denote by (v_s, v_s') the pair of vertices connected by the edge inserted in step $s \le t$. The edge was generated by step (2) of the sampling scheme if and only if one of the vertices has degree 0 in G_{s-1}, and hence if the indicator variable

$$B_s := \mathbb{1}\{\min\{\deg_{s-1}(v_s), \deg_{s-1}(v_s')\} = 0\}$$

takes value 1. The sequence's probability of occurrence under an **RW** model is then

$$C_{N_t,t}(\alpha) \prod_{s=2}^{t} \mathbb{P}\{\text{edge}_s = (v_s, v_s')\}^{1-B_s}, \tag{7}$$

where $C_{N_t,t}(\alpha)$ denotes the probability of $N_t - 2$ under a $\text{Binomial}(t-1, \alpha)$ distribution. The product's factors are given by Proposition 3.1: For a $\mathbf{RW}_{\text{SB}}(\alpha, \text{Poisson}_+(\lambda))$ model, for example, define

$$Q_s^{\lambda} := D_s^{-1/2}(\mathbb{I}_{N_s} - \mathbf{L}_s)\mathbf{K}_s^{\lambda} D_s^{1/2}. \tag{8}$$

Denote the relative degree of a vertex v in G_s by $\bar{d}_s(v) := \deg_s(v)/\sum_{v'\in\mathcal{V}G_s}\deg_s(v')$. Then $\hat{\alpha}_t$ as above and

$$\hat{\lambda}_t := \arg\max_{\lambda\geq 0}\prod_{s=2}^{t}\left([Q^{\lambda}_{s-1}]_{v_s v'_s}\bar{d}_{s-1}(v_s) + [Q^{\lambda}_{s-1}]_{v'_s v_s}\bar{d}_{s-1}(v'_s)\right)^{1-B_s} \tag{9}$$

are maximum likelihood estimators of the parameters α and λ. Neither estimate depends on the other, and $\hat{\lambda}_t$ can be computed by straightforward numerical optimization. Similarly, tractable ML estimators can be obtained in the negative binomial case, by substituting $\mathbf{K}^{r,p}$ for $\mathbf{K}^{\lambda}$ in (8). For a $\mathbf{RW}_{\mathrm{U}}(\alpha, \mathrm{Poisson}_+(\lambda))$ model, $\hat{\alpha}_t$ takes the same form as in (9); $\hat{\lambda}_t$ is obtained by replacing the relative degree $\bar{d}_{s-1}(v_s)$ with $1/N_{s-1}$.

For simple graphs, maximum likelihood estimation is still possible in principle, but more complicated since the effects of α and λ are no longer independent: In step (3') in Section 2, the probability of generating an additional vertex, which increments the count N_t, depends on the outcome of the random walk. The definition of step (3') is one of several possible ways to translate the multigraph case to simple graphs, but dependence of α and λ is not an artifact of a specific definition: Rather, it is due to the fact that a simple graph contains no observable evidence of a random walk connecting two previously connected vertices.

3.3. *Asymptotic degree properties.* A property of network models intensely studied in the theoretical literature is the behavior of vertex degrees as the graph grows large. For preferential attachment graphs, limiting degree distributions can be determined analytically [18, 43]. Our next results describe analogous properties for random walk models. The proofs use invariance of the degree-biased distribution under the operators $\mathbf{K}^{\lambda}$ and $\mathbf{K}^{r,p}$ to reduce to proof techniques developed for preferential attachment, despite the presence of the random walk.

Denote by $\mathcal{V}_{d,t}$ the subset of vertices in G_t with degree d, and let $m_{d,t} := |\mathcal{V}_{d,t}|$. The **degree distribution** is the normalized histogram vector $N_t^{-1}(m_{1,t}, m_{2,t}, \dots)$ of the vertex degrees in G_t. Let p_d be the probability that a vertex sampled uniformly at random has degree d. When the average probability (over vertices) of inserting a self-loop—i.e., of a random walk ending where it started—vanishes for large t, as is the case for $\lambda \to \infty$, it can be shown that as $t \to \infty$, $\mathbf{RW}_{\mathrm{SB}}$ random multigraphs have degree distribution of the form

$$p_d = \left(1 + \frac{\alpha}{2-\alpha}\right)\frac{\Gamma(d)\Gamma(2+\frac{\alpha}{2-\alpha})}{\Gamma(d+2+\frac{\alpha}{2-\alpha})}\,. \tag{10}$$

This distribution over $\mathbb{N}_+$ is known as the **Yule-Simon distribution** with parameter $\rho = 1 + \frac{\alpha}{2-\alpha}$ [e.g. 18]. For large d, p_d scales as a power law in d with exponent

$$\gamma := 2 + \frac{\alpha}{2-\alpha} = 1 + \rho \in (2, 3]\,. \tag{11}$$

For general λ that may have non-vanishing average self-loop insertion probabilities, simulations indicate the same behavior.

THEOREM 3.2 (Degree distribution). *Let a sequence of multigraphs $(G_1, G_2, \dots)$ have law $\mathbf{RW}_{\text{SB}}(\alpha, P)$, for some distribution P on $\mathbb{N}$, such that*

$$\frac{1}{m_{d,t}} \sum_{v \in \mathcal{V}_{d,t}} P(V_{end} = v | V_0 = v, G_t) = o(1) \qquad \text{as } t \to \infty \,. \tag{12}$$

Then the scaled number of vertices in G_t with degree d follows a power law in d with exponent γ as in (11). In particular, for all $d \in \mathbb{N}_+$,

$$\frac{m_{d,t}}{\alpha \cdot t} \longrightarrow p_d \qquad \text{in probability as } t \to \infty \,. \tag{13}$$

Condition (12) controls the number of self-loops in the graph, by requiring that the probability of a random walk ending where it starts vanishes for $t \to \infty$, separately for each degree d.

Denote by $\mathbf{deg}_t := (\deg_t(v_1), \deg_t(v_2), \dots)$ the **degree sequence** of a sequential random graph model, where v_j is the j-th vertex to appear in the graph sequence. For the preferential attachment model, the limiting degree sequence $\mathbf{deg}_\infty$ can be studied analytically, and has received considerable attention in the probability literature [see e.g. 18, 43]. The PA degree sequence is closely related to a Pólya urn, and like the limit of an urn, $\mathbf{deg}_\infty$ is itself a random variable: Informally, edges created early have sufficiently strong influence on later edges that randomness does not average out asymptotically. The joint law of $\mathbf{deg}_\infty$ can be obtained explicitly [35, 43], and determines the local weak limit [6]. In some cases, it also admits constructive representations: Using only sequences of independent elementary random variables, one can generate a random sequence whose joint distributions are identical to those of the limiting degree sequence [35, 27]. Such representations are closely related to "stick-breaking constructions" of random partitions [e.g. 39].

The next result similarly describes the limiting degree sequence $\mathbf{deg}_\infty$ of the $\mathbf{RW}_{\text{SB}}$ model. The relevant technical tool is to condition on the order in which edges occur: For a graph sequence $(G_1, G_2, \dots)$ be a graph sequence generated by the model, denote by S_j the time index at which vertex v_j is inserted in the model (that is, G_{S_j} is the first graph in sequence that contains v_j), and let $S_{1:r} := (S_1, \dots, S_r)$. Let Σ_β be a positive stable random variable with index β, i.e. characterized by the Laplace transform $\mathbb{E}[e^{-t\beta}] = e^{-t^\beta}$, and $f_\beta(s)$ its density. Define $\Sigma_{\beta,\theta}$, for $\theta > -\beta$, as a random variable with the polynomially tilted density $f_{\beta,\theta}(s) \propto s^{-\theta} f_\beta(s)$. The variable $M_{\beta,\theta} := \Sigma_{\beta,\theta}^{-\beta}$ is said to have **generalized Mittag-Leffler distribution** with parameters β and θ [39, 27].

THEOREM 3.3 (Degree sequence). *Let a sequence of multigraphs $(G_1, G_2, \dots)$ have law $\mathbf{RW}_{\text{SB}}(\alpha, P)$, for some distribution P on $\mathbb{N}$. Conditionally on $(S_1, S_2, \dots)$, the scaled degree sequence converges jointly to a random limit: For each $r \in \mathbb{N}_+$,*

$$t^{-1/\rho}(\deg_t(v_1), \deg_t(v_2), \dots, \deg_t(v_r)) \,\big|\, S_{1:r} \quad \xrightarrow{t\to\infty} \quad (\xi_1, \xi_2, \dots, \xi_r) \,\big|\, S_{1:r} \tag{14}$$

almost surely. Each limiting conditional law $\mathcal{L}(\xi_j \,|\, S_j)$, for $j \geq 1$, can be represented constructively: Let $M_j \sim \text{Mittag-Leffler}\left(\rho^{-1}, S_j - 1\right)$ and $B_j \sim \text{Beta}\left(1, \rho(S_j - 1)\right)$

be conditionally independent given S_j*. Then*

$$M_j \cdot B_j \big| S_j \stackrel{d}{=} \xi_j \big| S_j \ . \tag{15}$$

The marginal law $\mathcal{L}(\xi_j)$ *is uniquely characterized by the moments*

$$\begin{aligned}\mathbb{E}\big[\xi_j^k\big] &= \frac{\Gamma(k+1)\Gamma(j-1)}{\Gamma(j-1+\frac{k}{\rho})}\alpha^{\frac{k}{\rho}}\,{}_2F_1\big(1+\tfrac{k}{\rho},\tfrac{k}{\rho};j-1+\tfrac{k}{\rho};1-\alpha\big) \\ &= \Gamma(k+1)\alpha^{\frac{k}{\rho}}j^{-\frac{k}{\rho}}\cdot(1+O(j^{-1})) \qquad \text{as } j\to\infty\ ,\end{aligned} \tag{16}$$

where ${}_2F_1(a,b;c;z)$ *is the ordinary hypergeometric function.*

For any fixed T, one can fix an arbitrary enumeration of the edge set of G_T. The sequence $G_{1:T}$ then determines a (random) permutation Π of the edge set $\{1,\dots,T\}$, specifying the order in which edges where inserted, and $G_{1:T}$ can be represented equivalently as the pair (Π, G_T). We will find in Section 4 that Π also plays an important role in inference, as a latent variable. Π completely determines $(S_1,\dots,S_T)$ (though not vice versa), and the result above remains valid if one conditions on Π rather than $(S_{1:T})$. The simple constructive representation of the limiting distribution (15) seems only to exist marginally, not for the joint law $\mathcal{L}(\xi_1,\dots,\xi_r \,\big|\, S_{1:r})$, in contrast to the PA model, where a recursive analogue of (15) holds even jointly [27].

3.4. *Relation to preferential attachment models.* If the random walk step (3) in Algorithm 1 is omitted—that is, for $\alpha = 1$—the $\mathbf{RW}_{\text{SB}}$ model becomes reminiscent of the well-known preferential attachment model [4]. There are varying definitions of the PA model, but in essence, it inserts a new vertex into the graph at every step, and connects it to m vertices sampled from the size-biased distribution on the current vertex set; m is a model parameter.

A generalization of this model, due to Aiello, Chung, and Lu [1], is defined as follows: Fix $\alpha \in (0,1]$ and let G_1 be a graph with a two vertices connected by a single edge. In each step t, construct G_t from G_{t-1} as follows: Select a vertex V according to the size-biased distribution on G_{t-1}; with probability α, attach a new vertex; otherwise, connect V to a vertex V' sampled independently from the same size-biased distribution. This model is denoted $\mathbf{ACL}(\alpha)$ in the following. It can be regarded as the natural extension of the Yule-Simon model from sequences to graphs.

PROPOSITION 3.4. *The limit in distribution* $\mathbf{RW}_{\text{SB}}(\alpha,\infty) := \lim_{\lambda\to\infty}\mathbf{RW}_{\text{SB}}(\alpha,\lambda)$ *exists for every* α*, and* $\mathbf{RW}_{\text{SB}}(\alpha,\infty) = \mathbf{ACL}(\alpha)$ *if both models start with the same seed graph.*

4. Particle methods for sequential network models. This section develops Markov chain sampling methods for a commonly encountered problem: A sequential network model is assumed, but only the final graph G_T generated by the model is observed, as opposed to the entire history $G_{1:T} := (G_1,\dots,G_T)$ of the

generative process. The methods developed here assume a Bayesian setup, with a prior distribution $\mathcal{L}(\theta)$ on the model parameter θ, where $\mathcal{L}(\,\bullet\,)$ generically denotes the law of a random variable. They sample the posterior distribution $\mathcal{L}(\theta|G_T)$, and are applicable to sequential network models satisfying two properties:

(P1) The sequence $(G_1, \ldots, G_T)$ of graphs generated by the models forms a Markov chain on the set of finite graphs, that is, $G_{t+1} \perp\!\!\!\perp_{G_t} (G_1, \ldots, G_{t-1})$ for $t < T$.
(P2) The sequence is strictly increasing, in the sense that $G_t \subsetneq G_{t+1}$ almost surely.

These hold for the random walk models, but also for many other network formation models, such as preferential attachment graphs, fitness models, and vertex copying models [e.g. 37, 22]. Despite the attention these models have received in the literature, little work exists on inference (see Section 4.5 for references).

If only a single graph G_T is observed, inference requires that the unobserved history is integrated out. The result is a likelihood of the form

$$p_\theta(G_T \,|\, G_1) = \int p_\theta(G_T, G_{2:(T-1)} \,|\, G_1)\; dG_{2:(T-1)} \;, \tag{17}$$

where θ is the vector of model parameters. Since the variables G_t take values in large combinatorial sets, the integral amounts to a combinatorial sum that is typically intractable. The strategy is to approximate the integral with a sampler that imputes the unobserved graph sequence, noting that only valid sequences which lead to G_T will have non-zero likelihood (17). The sequential nature of models makes SMC and particle methods the tools of choice.

4.1. *Basic SMC.* We briefly recall SMC algorithms: The canonical application is a state space model. Observed is a sequence $X_{1:T} = (X_1, \ldots, X_T)$. The model explains the sequence using an unobserved sequence of latent states $Z_{1:T}$, and defines three quantities:

- A Markov kernel $q^t_\theta(\,\bullet \,|\, Z_{t-1})$ that models transitions between latent states.
- An emission distribution p^t_θ that explains each observation as $X_t \sim p^t_\theta(\,\bullet \,|\, Z_t)$.
- A vector θ collecting the model parameters.

In the simplest case, θ is fixed. The inference target is then the posterior distribution $\mathcal{L}_\theta(Z_{1:T} \,|\, X_{1:T})$ of the latent state sequence.

A SMC algorithm generates some number $N \in \mathbb{N}$ of state sequences $Z^1_{1:T}, \ldots, Z^N_{1:T}$, and then approximates the posterior as a sample average over these sequences, weighted by their respective likelihoods. The imputed states $Z^i_{1:t}$ are called **particles**. Since the sequence of latent states is Markov, particles can be generated sequentially as $Z_t \sim q^t_\theta(Z_t|Z_{t-1})$. In cases where sampling from q^t_θ is not tractable, q^t_θ it is additionally approximated by a simpler proposal kernel r^t_θ. The likelihood, and the accuracy of approximation of q^t_θ by r^t_θ, are taken into account by computing normalized weights

$$w^i_t := \frac{\tilde{w}^i_t}{\sum_{j=1}^N \tilde{w}^j_t} \qquad \text{where} \qquad \tilde{w}^i_t = p^t_\theta(X_t|Z^i_t) \cdot \frac{q^t_\theta(Z^i_t \,|\, Z^i_{t-1})}{r^t_\theta(Z^i_t \,|\, Z^i_{t-1})} \;. \tag{18}$$

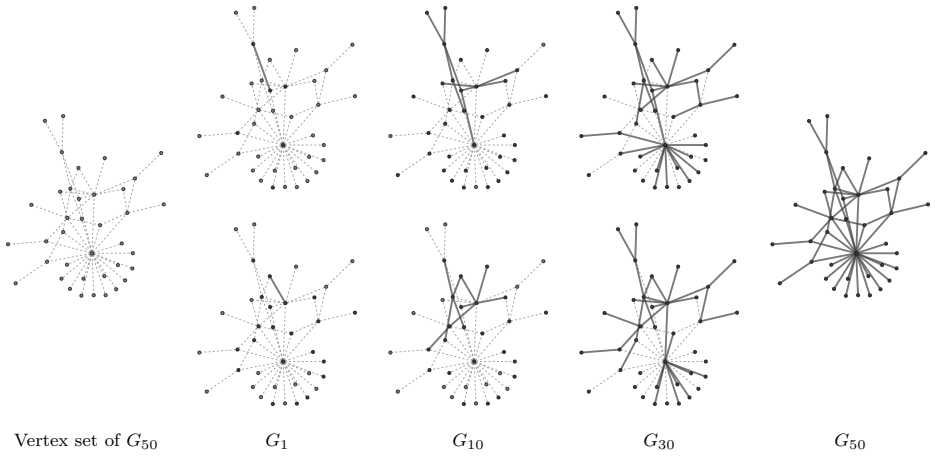

FIG 2. *Two graph bridges generated by Algorithm 2: A graph G_{50} is drawn from a $\mathbf{RW}_{\mathrm{U}}(\alpha, P)$ model, and two graph bridges $G_{1:50}$ are sampled conditionally on the fixed input G_{50}. Shown are the graphs G_1, G_{10} and G_{30} of each bridge.*

In step t, SMC generates particles Z_t^i by first resampling from the previous particles $(Z_{1:t-1}^i)_i$ with probability proportional to their weights, i.e. by drawing

$$A_t^i \sim \mathrm{Multinomial}(N, (w_{t-1}^i)_{i=1}^N) \; .$$

The next set $(Z_{1:t}^i)_i$ of particles is then generated by drawing $Z_t^i \sim r_\theta^t(\,\bullet \mid Z_{t-1}^{A_t^i})$ and setting $Z_{1:t}^i = (Z_{1:t-1}^{A_t^i}, Z_t^i)$. Resampling a final time after the T-th step yields a complete array $(Z_{1:T}^i)_i$, with which the posterior is approximated as the average

$$\mathcal{L}_\theta(dz_{1:T} \,|\, X_{1:T}) \approx \frac{1}{N} \sum_{i=1}^{N} \delta_{Z_{1:T}^i}(dz_{1:T}) \; . \tag{19}$$

Alternatively, the final resampling step can be omitted, in which case the posterior is approximated with a weighted average. Either approximation is asymptotically unbiased as $N \to \infty$. See Doucet and Johansen [17] for a thorough review.

4.2. *SMC algorithms for graph bridges.* Consider a sequential network model satisfying properties (P1) and (P2) above, with model parameters collected in a vector θ. For now, θ is fixed, and the objective is to reconstruct the graph sequence $G_{1:T}$ from its observed final graph G_T and a fixed initial graph G_1, i.e. the relevant posterior distribution is $\mathcal{L}_\theta(G_{1:T} \,|\, G_T, G_1)$. By the Markov property (P1), each step in the process is governed by a Markov kernel

$$q_\theta^t(g \,|\, g') := P(G_t = g \,|\, G_{t-1} = g') \; .$$

The task of a sampler is hence to impute the conditional sequence $G_{2:(T-1)} \,|\, G_1, G_T$, which is a stochastic process conditioned on its initial and terminal point, also

known as a **bridge**. Compared to the SMC sampling algorithm for state space models, the unobserved sequence $G_2, \dots, G_{T-1}$ takes the role of the hidden states $z_{1:T}$, whereas G_1 and the single observed graph G_T replaces the observed sequence $x_{1:T}$. The emission likelihood $p_\theta^t(x_t|z_t)$ is replaced by a bridge likelihood, of a graph G_t given G_1 and G_T.

The relevant likelihood functions, however, are themselves intractable: If g_s is a fixed graph at index s in the sequence, the probability of observing g_t at a later point $t > s$ is

$$p_\theta^{t,s}(g_t|g_s) := \int \Big(\prod_{i=s}^{t-1} q_\theta^{i+1}(g_{i+1}|g_i)\Big) dg_{s+1} \cdots dg_{t-1} \qquad \text{whenever } t > s \ .$$

For a candidate graph G_t generated by the sampler, the likelihood under observation G_T is the **bridge likelihood**

$$L_\theta^t(G_t) := p_\theta^{T,t}(G_T|G_t) \ .$$

The bridge likelihood is intractable unless $t = T - 1$; indeed, for $t = 1$, this is precisely the integral (17) that we set out to approximate.

A sequence $G_{1:T}$ satisfying (P2) is equivalent to an enumeration of the T edges of G_T in order of occurrence. Let π be a permutation of $\{1, \dots, T\}$, i.e. an ordered list containing each element of the set exactly once, and

$$\pi_t(G_T) := \text{graph obtained by deleting all edges of } G_T \text{ not listed in } \pi_t, \text{ and any resulting isolated vertices.}$$

A random sequence $G_{1:T}$ can then be specified as a pair (Π, G_T), for a *random* permutation Π of the edges in G_T, with $G_t = \Pi_t(G_T)$. Given G_T, not every permutation π is a valid candidate value for the unobserved order: π must describe a sequence of steps of non-zero probability from the initial graph to G_T, and we hence have to require

$$L_\theta^t(\pi_t(G_T)) > 0 \qquad \text{for all } t \leq T \ . \tag{20}$$

If so, we call π **feasible** for G_T. For G_T given and $G_t = \Pi_t(G_T)$, we use G_t and Π_t interchangeably. The target distribution of the SMC bridge sampler at step t is then

$$\gamma_\theta(t) := L_\theta^t(\Pi_t)P(\Pi_t) = L_\theta^t(\Pi_t) \prod_{s=1}^{t-1} q_\theta^{s+1}(\Pi_{s+1} \,|\, \Pi_s) \ , \tag{21}$$

which satisfies the recursion

$$\gamma_\theta(t) = \begin{cases} \frac{L_\theta^t(\Pi_t)}{L_\theta^{t-1}(\Pi_{t-1})} q_\theta^t(\Pi_t|\Pi_{t-1})\gamma_\theta(t-1) & \text{if } \Pi_t \text{ is feasible} \\ 0 & \text{otherwise} \end{cases} .$$

The intractable part is the bridge likelihood ratio $L_\theta^t / L_\theta^{t-1}$. Define

$$h_t(\Pi_t) := \begin{cases} \mathbb{1}\{L_\theta^t(\Pi_t) > 0\} & \text{for all } t < T-1 \\ L_\theta^t(\Pi_t) & \text{for } t = T-1 \end{cases} .$$

As Proposition 4.1 below shows, the apparently crude approximation $L_\theta^t \approx h_t$ still produces asymptotically unbiased samples from the posterior $\mathcal{L}_\theta(G_{1:T} \,|\, G_T, G_1)$; this is based on methodology developed in [16] for bridges in continuous state spaces. If Π_t is feasible, substituting h_t into (21) yields the surrogate recursion

$$\gamma_\theta(t) = \frac{q_\theta^t(\Pi_t \,|\, \Pi_{t-1})}{r_\theta^t(\Pi_t \,|\, \Pi_{t-1})}\gamma_\theta(t-1) \qquad \text{for} \quad t < T-1 \,.$$

The SMC proposal kernel r_θ^t is hence chosen as the truncation of q_θ^t to feasible permutations,

$$r_\theta^t(\Pi_t \,|\, \Pi_{t-1}) := \frac{\mathbb{1}\{L_\theta^t(\Pi_t) > 0\} q_\theta^t(\Pi_t \,|\, \Pi_{t-1})}{\tau_\theta^t(\Pi_{t-1})} \quad \text{with} \quad \tau_\theta^t := \sum_{\pi_t : L_\theta^t(\pi_t) > 0} q_\theta^t(\pi_t \,|\, \Pi_{t-1}) \,. \tag{22}$$

For particles $G_t^i = \Pi_t^i(G_T)$, the unnormalized SMC weights (18) are

$$\tilde{w}_t^i = \begin{cases} \tau_\theta^t(\Pi_{t-1}^i) & \text{if } t < T-1 \\ q_\theta^t(\Pi \,|\, \Pi_{T-1}^i)\, \tau_\theta^{T-1}(\Pi_{T-2}^i) & \text{if } t = T-1 \end{cases} . \tag{23}$$

The sampling algorithm then generates a graph bridge as follows:

Algorithm 2 (Bridge sampling)**.**

- Initialize $G_1^i := G_1$, $\Pi_1^i \sim \text{Uniform}\{1, \dots, T\}$, and $w_1^i := 1/N$ for each $i \leq N$.
- For $t = 2, \dots, T-1$, iterate:
  - Resample indices $A_t^i \sim \text{Multinomial}(N, (w_{t-1}^i)_i)$.
  - Draw $\Pi_t^i \sim r_\theta^t(\,\bullet \,|\, \Pi_{t-1}^{A_t^i})$ as in (22) for each i.
  - Compute weights as in (23) and normalize to obtain w_t^i.
- Resample N complete sequences $G_{1:T}^i = \Pi^i \sim \text{Multinomial}(N, (w_{T-1}^i)_i)$.

See Figure 2 for an illustration. Computing $h_t(\Pi_t)$ and $\tau_\theta^t(\Pi_{t-1})$ is simplified by the constraints on Π: Given Π_{t-1}, the requirement that Π_t must again be a restriction of a permutation Π implies $\Pi_t = (\Pi_{t-1}, s_t)$, for some $s_t \in \{1, \dots, T\} \setminus \Pi_{t-1}$. Within this set, the s_t for which Π_t is feasible are simply those edges in G_T connected to $\Pi_{t-1}(G_T)$.

PROPOSITION 4.1. *Let q_θ^t, for $t = 1, \dots, T$, be the Markov kernels defining a sequential network model that satisfies conditions (P1) and (P2). Given an observation G_T, Algorithm 2 produces samples that are asymptotically unbiased as $N \to \infty$. That is, for any bounded function f on graph sequences,*

$$\frac{1}{N}\sum_{i=1}^N f(G_{1:T}^i) \xrightarrow{p} \mathbb{E}[f(G_{1:T}) \,|\, G_T, G_1] \qquad \text{as } N \to \infty \,, \tag{24}$$

where the expectation is evaluated with respect to the model posterior $\mathcal{L}_\theta(G_{1:T} \,|\, G_T, G_1)$.

REMARK. If estimates exhibit high variance, it is straightforward to modify Algorithm 2 to use adaptive resampling [see 16], and to replace multinomial resampling above by residual or stratified resampling [e.g. 17]. If a given model admits a more bespoke approximation to the bridge likelihood, this approximation can be substituted for h_t, following [16]. In this case, some of the equations above require (elementary) modifications.

4.3. *Parameter inference.* Algorithm 2 generates a history of a graph under a model with fixed parameter vector θ. For parameter inference, the parameters are treated as a random variable Θ, with prior distribution $\mathcal{L}(\Theta)$, and the task is to generate samples from the joint posterior $\mathcal{L}(\Theta, G_{1:T}|G_1, G_T)$. The sample space of the SMC sampler above is thus extended by the domain of Θ. The bridge likelihood $L^1_\Theta(G_1)$ is a marginal likelihood that naturally leads to particle MCMC methods [3], which use SMC to compute a consistent estimate $\hat{L}^1_\Theta$. Substituting into the Metropolis-Hastings acceptance ratio of a proposal $\tilde{\Theta}$ from a (yet to be specified) proposal distribution $\tilde{q}$ yields

$$\mathbb{P}\{\text{ accept } \tilde{\Theta}\,\} = \frac{\hat{L}^1_{\tilde{\Theta}} \cdot \pi(\tilde{\Theta})}{\hat{L}^1_{\Theta} \cdot \pi(\Theta)} \cdot \frac{\tilde{q}(\Theta \,|\, \tilde{\Theta})}{\tilde{q}(\tilde{\Theta} \,|\, \Theta)} \,. \tag{25}$$

By Proposition 4.1, a consistent estimate $\hat{L}^1_\Theta$ of the bridge likelihood can be computed from the final set of unnormalized weights generated by Algorithm 2 as

$$\hat{L}^1_\Theta := \frac{1}{N} \sum_{i=1}^{N} \tilde{w}^i_{T-1} \,. \tag{26}$$

Using (25) and (26), Algorithm 3 defines a particle marginal Metropolis-Hastings (PMMH) sampler.

Algorithm 3.

- Initialize $\Theta^0 \sim \pi$.
- For $j = 1, \ldots, J$ iterate:
  1. Draw candidate a value $\tilde{\Theta} \sim \tilde{q}(\,\bullet \,|\, \Theta^{j-1})$.
  2. Run Algorithm 2 with parameter $\tilde{\Theta}$ to compute $\hat{L}^1_{\tilde{\Theta}}$ as in (26).
  3. Accept $\tilde{\Theta}$ with probability (25) and set $\Theta^j := \tilde{\Theta}$; else set $\Theta^j := \Theta^{j-1}$.
  4. If $\tilde{\Theta}$ is accepted, select a single graph sequence $G^j_{1:T}$ by resampling from the particles output by Algorithm 2; else set $G^j_{1:T} = G^{j-1}_{1:T}$.
- Output the sequence $(\Theta^1, G^1_{1:T}), \ldots, (\Theta^J, G^J_{1:T})$.

The algorithm asymptotically samples the joint posterior $\mathcal{L}(\Theta, G_{1:T} \,|\, G_1, G_T)$, or the marginal posterior $\mathcal{L}(\Theta \,|\, G_1, G_T)$ if step (d) is omitted:

PROPOSITION 4.2. *If the proposal density $\tilde{q}(\,\bullet\,|\,\bullet\,)$ is chosen such that the Metropolis-Hastings sampler defined by (25) is irreducible and aperiodic, Algorithm 3 is a PMMH sampler. The marginal distributions $\mathcal{L}(\Theta^j, G^j_{1:T})$ of its output sequence satisfy*

$$\|\mathcal{L}(\Theta^j, G^j_{1:T}) - \pi(\,\bullet\,|\,G_T, G_1)\|_{\text{TV}} \xrightarrow{j\to\infty} 0\ . \tag{27}$$

This is true regardless of the sample size N generated by Algorithm 2 in step (b).

Although the result holds asymptotically independently of N, a larger value of N will generally speed up convergence and reduce variance.

4.4. *Particle Gibbs for random walk models.* The computational cost of Algorithm 3 stems mostly from two sources: Each rejection in (c) requires an additional execution of steps (a) and (b), and the cost of each such execution may be high. Details depend on the network model, but in most cases, (b) is the most expensive step. Rejections can be addressed by turning the MH algorithm into a Gibbs sampler, which eliminates the acceptance step, and does not require the choice of a proposal kernel $\tilde{q}$. The resulting algorithm constitutes a particle Gibbs (PG) sampler [3]. Such an algorithm is described below, now specifically for the random walk model. The algorithm uses two sequences of auxiliary variables B_t and K_t, each corresponding to one model parameter. At each step, some of these variables, as well as the model parameters themselves, can be integrated out, which simplifies sampling significantly.

For a $\mathbf{RW}(\alpha, \text{Poisson}_+(\lambda))$ model, the parameter takes the form $\Theta = (\alpha, \lambda)$, and we fix a beta distribution $P_{[\alpha]}$ as prior for α, and a gamma prior $P_{[\lambda]}$ for λ. To sample a sequence $G_1, \ldots, G_T$ from the model given a pair (α, λ), one can generate two i.i.d. sequences, $\mathbf{B} = (B_1, \ldots, B_T)$ of Bernoulli(α) variables, and $\mathbf{K} = (K_1, \ldots, K_T)$ of Poisson$_+(\lambda)$ variables. In step t of Algorithm 1, a new edge is inserted if $B_t = 1$; otherwise, two vertices are connected by random walk of length K_t. Since $G_{1:T}$, $\mathbf{B}$ and $\mathbf{K}$ are dependent random variables, the kernel $q^t_{\alpha,\lambda}$ in Algorithm 2 depends explicitly on B_t and K_t, a function of G_{t-1}, B_t and K_t. As shown in Appendix B, the entries B_t and K_t, along with the model parameters α and λ, can be integrated out of the kernel, such that Algorithm 2 samples

$$G_{t+1} \sim q^{t+1}_{\alpha,\lambda}(\,\bullet\,|\,G_t, \mathbf{B}_{-t}, \mathbf{K}_{-t}) \qquad \text{where} \qquad \mathbf{B}_{-t} := (B_1, \ldots, B_{t-1}, B_{t+1}, \ldots, B_T)\ ,$$

and $\mathbf{K}_{-t}$ is defined analogously. $\mathbf{B}$ and $\mathbf{K}$ can hence be sampled separately from the SMC steps, rather than inside Algorithm 2, and improves exploration of the sample space.

In its jth iteration, the algorithm updates $\mathbf{B}$ and $\mathbf{K}$ by looping over the indices $t \le T$ of $G_{1:T}$. Since Gibbs samplers condition on every update immediately, vectors maintained by the sampler are of the form $\mathbf{B}^j_{-t} := (B^{j+1}_1, \ldots, B^{j+1}_{t-1}, B^j_t, \ldots, B^j_T)$, with $\mathbf{K}^j_{-t}$ again defined analogously. Marginalizing out α, K_t, and λ yields the posterior predictive distribution of B^{j+1}_t,

$$\mathcal{L}(B^{j+1}_t\,|\,\mathbf{B}^j_{-t}, \mathbf{K}^j_{-t}, G^j_{(t-1):t}) = \int q^t_{\alpha,\lambda}(G^j_t\,|\,G^j_{t-1}, \mathbf{K}^j_{-t}, \bullet)\text{Bernoulli}(\,\bullet\,|\,\alpha)P_{[\alpha]}(d\alpha\,|\,\mathbf{B}^j_{-t})\ ,$$

where $P_{[\alpha]}(\alpha \,|\, \mathbf{B}^j_{-t})$ is the posterior distribution of α given $\mathbf{B}^j_{-t}$. Similarly, the predictive distribution of K_t^{j+1} marginalizes out λ, B_t^{j+1}, and α, and both distributions can be obtained in closed form, despite their dependence on the graph bridge $G^j_{1:T}$ (see Appendix B for details). In the algorithm below, we abuse notation and let the index $j = 0$ refer to the predictive distribution under the prior, i.e. we write $\mathcal{L}(\mathbf{B}^1|\mathbf{B}^0, G^0_{1:T}) := \int \mathcal{L}(\mathbf{B}^1 \,|\, \alpha) P_{[\alpha]}(d\alpha)$, and similarly for $\mathcal{L}(\mathbf{K}^1 \,|\, \mathbf{K}^0, G^0_{1:T})$.

Algorithm 4 (Particle Gibbs sampling for $\mathbf{RW}(\alpha, \text{Poisson}_+(\lambda))$ models.)**.**

- For $j = 0, \dots, J$:
  1. Draw $(\mathbf{B}^{j+1}, \mathbf{K}^{j+1}) \sim \mathcal{L}(\mathbf{B}^{j+1}, \mathbf{K}^{j+1} \,|\, \mathbf{B}^j, \mathbf{K}^j, G^j_{1:T})$.
  2. Using Algorithm 2, with $q^t_{\alpha,\lambda}(\,\bullet \,|\, G^j_{t-1}, \mathbf{B}^{j+1}_{-t}, \mathbf{K}^{j+1}_{-t})$ and $N \geq 2$, update $G^{j+1}_{1:T}$ and $a^{j+1}_{2:T}$. At each step t, substitute the previous iteration's bridge $G^j_{1:t}$ for the particle with index a^j_t (see below).
  3. Draw $\alpha^{j+1} \,|\, \mathbf{B}^{j+1}$ and $\lambda^{j+1} \,|\, \mathbf{K}^{j+1}$ from their conjugate posteriors.
- Output the sequence $(G^1_{1:T}, \alpha^1, \lambda^1), \dots, (G^J_{1:T}, \alpha^J, \lambda^J)$.

This algorithm is used for inference in the random walk model in all experiments reported in Section 5. It constitutes a blocked Gibbs sampler with blocks (α, λ), $(\mathbf{B}, \mathbf{K})$, and $G_{1:T}$. Due to the marginalization described above, the parameter values α^j and λ^j generated in (c) are not used in the execution of the sampler; they only serve as output. Step (b), which seeds Algorithm 2 with a bridge $G^j_{1:T}$ generated during the previous iteration, is a **conditional SMC** step that biases the sampler towards $G^j_{1:T}$ and is necessary to make Algorithm 4 a valid Gibbs sampler. Although the resampled indices a_t have no effect on the sampled sequence $G_{1:T}$, the Gibbs sampler is defined on an augmented space that includes the random variables generated during the SMC steps, and so requires that we include them in the conditional SMC update. See [3] for more on conditional SMC.

COROLLARY 4.2.1. *Algorithm 4 is a valid particle Gibbs sampler. For any $N \geq 2$, it generates a sequence $(G^j_{1:T}, \alpha^j, \lambda^j)_j$ whose marginal laws converges as*

$$\|\mathcal{L}(G^j_{1:T}, \alpha^j, \lambda^j) - \mathcal{L}(G_{1:T}, \alpha, \lambda \,|\, G_T, G_1)\|_{\text{TV}} \xrightarrow{j\to\infty} 0\,. \tag{28}$$

to the model posterior $\mathcal{L}(G_{1:T}, \alpha, \lambda \,|\, G_T, G_1)$.

4.5. *Related work.* Existing work on the use of sampling algorithms for inference in sequential network models mainly follows one of two general approaches: Importance sampling [46], or approximate Bayesian computation (ABC) schemes [42]. ABC relies on heuristic approximations of the likelihood, and therefore is not guaranteed to sample from the correct target distribution. The basic approach used above—conduct inference by imputing graph sequences generated by a suitable sampler—seems to be due to Wiuf *et al.* [46], based on earlier work by Griffiths and

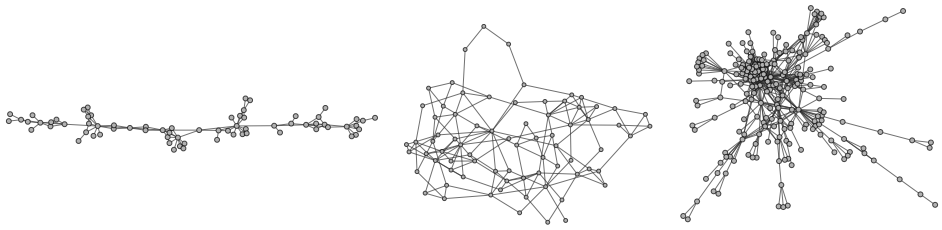

FIG 3. *Network data examples: (i) the largest connected component of the NIPS co-authorship network, 2002-03 [21]; (ii) San Juan Sur family ties [33]; (iii) protein-protein interactome [11].*

Tavare [23, 24] on ancestral inference for coalescent models. The work of Wang *et al.* [45] is related to ours in that it employs particle MCMC, for a particular sequential model. All these methods are, in short, applicable if the sequential model in question is very simple. Otherwise, they suffer (1) from the high variance of estimates that is a hallmark of importance sampling [17], and (2) infeasible computational cost. The algorithm in [45], for example, samples backwards through the sequence generated by the model, and for each reverse step $G_t \to G_{t-1}$ requires computing the probability of every possible forward step $G'_{t-1} \to G_t$ under the given model; even for the (still rather simple) random walk model, that is no longer practical.

A separate body of work applies SMC to a related problem, inference in probabilistic graphical models [e.g. 36]. Here, models lack a sequential structure, which is addressed algorithmically by constructing the dependence structure clique by clique via artificial intermediate distributions; not surprisingly, the design of good intermediate distributions turns out to be crucial. In a sequential graph model, these intermediate distributions are given by the model.

5. Experimental evaluation. This section evaluates properties of the model on real-world data and compares its performance to other network models. We also discuss the interpretation of the random walk parameter as a length scale, and of the latent order Π as a measure of vertex centrality.

We consider three network datasets, shown in Figure 3, that exhibit a range of characteristics. The first is the largest connected component (LCC) of the NIPS co-authorship network in 2002-03, extracted from the data used in [21]. As shown in Figure 3, it has a global chain structure connecting highly localized communities. The second dataset represents ties between families in San Juan Sur (SJS), a community in rural Costa Rica [33, 5], and is chosen here as an example of a network with small diameter. The third is the protein-protein interactome (PPI) of [11]. The PPI exhibits features such as chains, pendants, and heterogeneously distributed hubs. Basic summary statistics are as follows:

Dataset	*Vertices*	*Edges*	*Clustering coeff.*	*Diameter*	*L_2-mixing time of r.w.*
NIPS	70	114	0.70	14	≥ 62.4
SJS	75	155	0.32	7	≥ 6.1
PPI	230	695	0.32	11	≥ 9.2

None of these data sets is particularly large: Sampler diagnostics show graphs of this size suffice to reliably recover model parameters under the random walk model.

Using Algorithm 4, the posterior distributions of α and λ for both the uniform and the size-biased $\mathbf{RW}(\alpha, \text{Poisson}_+(\lambda))$ model can be sampled. Kernel-smoothed posterior distributions under all three datasets are shown in Figure 6. Although posteriors in the uniform and size-biased case are very similar, these models generate graphs with different characteristics for identical parameter values (see Section 5.3).

5.1. *Sampler diagnostics (synthetic data).* To assess sampler performance, we test the sampler's ability to recover parameter values for graphs generated by the model. For each such experiment, a single graph G_T with $T = 250$ edges is generated for fixed values of α and λ. The joint posterior distribution of α and λ given G_T is then estimated using Algorithm 4, run with a moderate number ($N = 100$) of particles. Regardless of the input parameter value, a uniform prior is chosen for α, and a Gamma$(1, 4)$ prior for λ. For the sake of brevity, we report results only for the $\mathbf{RW}_{\text{U}}(\alpha, \text{Poisson}_+(\lambda))$ model, on simple graphs (which pose a harder challenge for inference than multigraphs, since edge multiplicities provide additional information about the history of the graph). Figure 4 shows example traces of the samplers and the corresponding autocorrelation functions. The samplers mix and converge quickly, and as one would expect, only λ displays any noticeable autocorrelation. Posteriors for various input parameters are shown in Figure 5. Clearly, recovery of model parameters from a single graph is possible. The effect of the model's modification to generate simple graphs—the use of step (3') rather than (3) in Algorithm 1—is apparent for $\lambda = 2$: For small values of λ, a significant proportion of random walks ends at their starting point, resulting in the insertion of a new vertex. Since the model otherwise inserts new vertices as an effect of α, the two parameters are correlated in the posterior, which is clearly visible in Figure 5, especially for intermediate values of α.

5.2. *Length scale.* The distance of vertices that can form connections under the model is governed by the parameter λ of the random walk, which can hence be interpreted as a form of length-scale. This scale can be compared to the mixing time of a random walk on the observed graph (listed, in L_2 norm, in the table above for each data set): If λ is significantly smaller, the placement of edges inserted by the random walk strongly depends on the graph structure; if λ is large relative to mixing time, dependence between edges is weak. Based on Figure 6, we observe the following:

- Concentration of the λ-posterior on small values in $[1, 2]$ for the NIPS data thus indicates predominantly short-range dependence in the data and, since the mixing time is much larger, strong dependence between edges. Concentra-

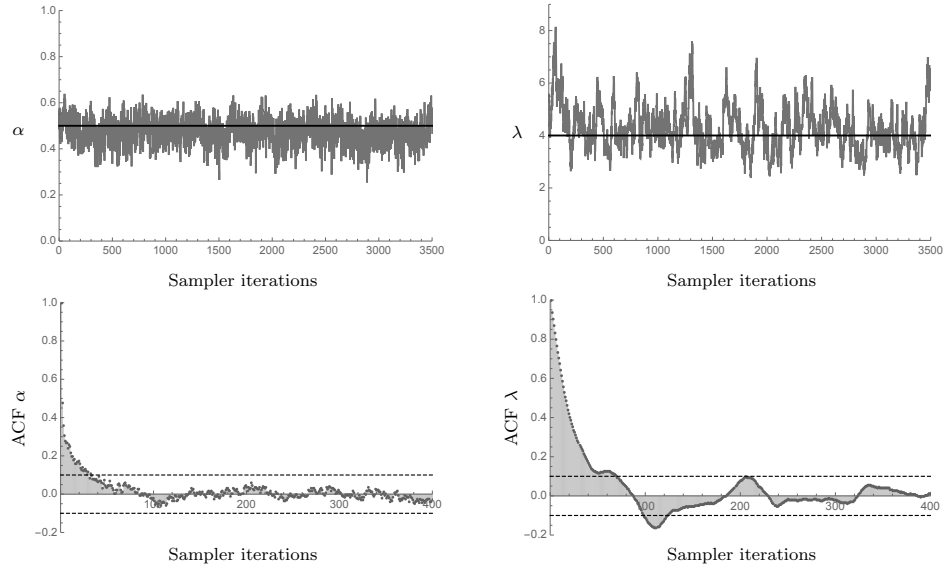

FIG 4. *Top: Posterior sample traces for a PG sampler fit to a synthetic graph G_T generated with $\alpha = 0.5$ and $\lambda = 4$ (dashed lines). Bottom: The corresponding autocorrelation functions.*

tion of the α-posterior near the origin means connections are mainly formed through existing connections.

- In the SJS network, the posterior peaks near $\lambda = 6$, and hence near the lower bound on the mixing time, with α again small. Thus, the principal mechanism for inserting edges under the model is again the random walk, but a random walk of typical length has almost mixed. Hence, the local connections it creates do not strongly influence each other.
- The PPI network exhibits an intermediate scale of dependence, with most posterior mass in the range $\lambda \in [3, 5]$. Although structures like pendants and chains indicate strong local dependence, there are also denser, more homogeneously connected regions that indicate weaker dependence with longer range.

Narratives explaining these effects are not hard to come by—for example, short-range dependence in the NIPS data indicate new collaborations are often formed through existing co-authors; the weak dependence between edges in the SJS data suggests connections can be formed by some process external to the observed graph, as may be expected for a community concentrated in a small geographic area—but warrant caution, as all data storytelling does.

5.3. *Model fitness.* Assessing model fitness on network data is difficult: For other types of data, cross validation is often the tool of choice, and indeed, cross-validated link prediction (where links and non-links are deleted uniformly and independently at random) is widely used for model evaluation in the machine learning literature. Even if link prediction is considered the relevant statistic of interest,

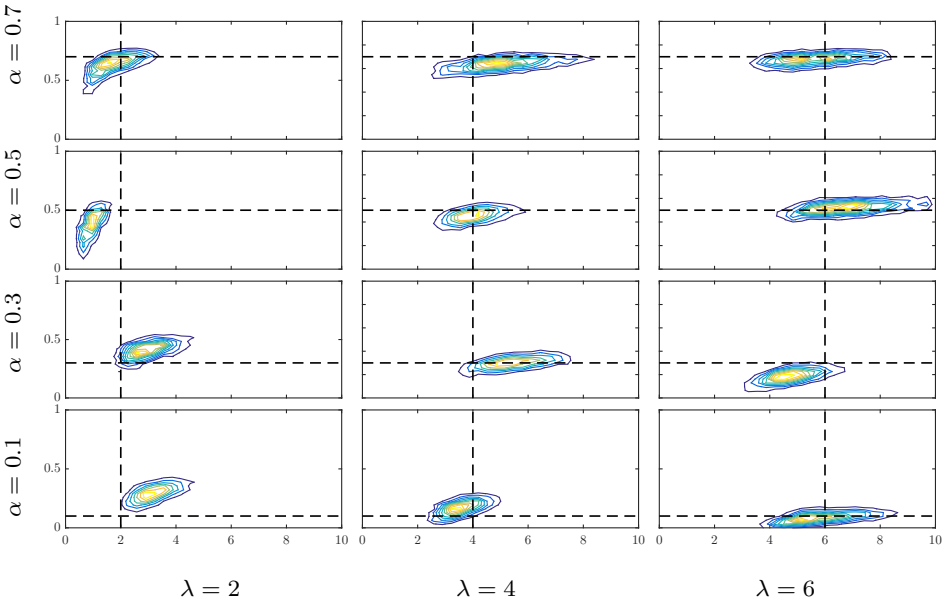

FIG 5. *Joint posterior distributions, given graphs generated from an* $\mathbf{RW}_{\mathrm{U}}$ *model for different parameter settings.* α *is on the vertical axis,* λ *on the horizontal axis, and generating parameter values are shown by dashed lines.*

however, cross-validating network data requires subsampling the observed network. Any choice of such a subsampling algorithm implies strong assumptions on how the data was generated; it may also favor one model over another.

We therefore use a protocol developed by Hunter *et al.* [26], who compare models to data by fixing a set of network statistics. The fitted model is evaluated by comparing the chosen statistics of the data to those of networks simulated from the model. The selected statistics specify which properties are considered important in assessing fit. To capture a range of structures, the following count statistics are proposed in [26], which we also adopt here:

- Normalized degree statistics (ND), d_k, the number of vertices of degree k, divided by the total number of vertices.
- Normalized edgewise shared partner statistics (NESP), χ_k, the number of unordered pairs $\{i, j\}$ such that i and j are connected and have exactly k common neighbors, divided by the total number of edges.
- Normalized pairwise geodesic statistics (NPG), γ_k, the number of unordered pairs $\{i, j\}$ with distance k in the graph, divided by the number of dyads.

In a Bayesian setting, executing the protocol amounts to performing posterior predictive checks [10, 20] via the following procedure:

(1) Sample the model parameters from the posterior, $\theta \sim \pi(\theta|G_T)$.

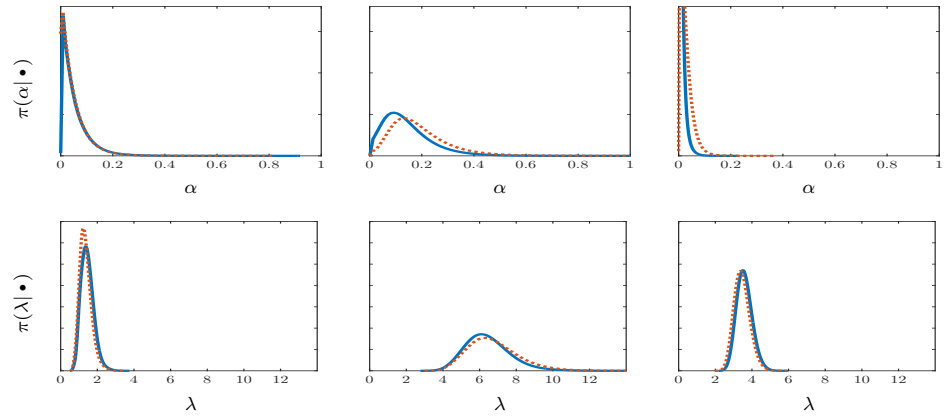

FIG 6. *Kernel-smoothed estimates of the posterior distributions of α and λ, under the models $\mathbf{RW}_{\mathrm{U}}$ (blue/solid) and $\mathbf{RW}_{\mathrm{SB}}$ (orange/dotted). Left column: NIPS data. Middle: SJS. Right: PPI. Posteriors are based on 1000 samples each (lag 40, after 1000 burn-in iterations; 100 samples each are drawn from 10 chains).*

(2) Simulate a graph of the same size as the data, $G^{(s)} \sim P(G|\theta)$.
(3) Calculate the statistic(s) of interest for the simulated graph, $f(G^{(s)})$.

The value $f(G^{(s)})$ is then a sample from the posterior predictive distribution (PPD) of the statistic f under the model. If the PPD places little mass on the observed value of the statistic, this indicates the model does not explain those properties of the data measured by the given statistic.

For each of the datasets and statistics above, PPDs are estimated for the following models: The Erdös-Rényi model (ER); the infinite relational model (IRM) with a Chinese Restaurant Process prior on the number of blocks [28, 48]; the infinite edge partition model with an underlying hierarchical gamma process (EPM) [49]; the ACL model described in Section 3.4; and the $\mathbf{RW}_{\mathrm{U}}$ and $\mathbf{RW}_{\mathrm{SB}}$ models. The following table lists the total variation distance between the empirical distribution of each statistic on the observed graph and on graphs generated from the respective models. Standard errors are computed over 1000 samples. Smaller values indicate a better fit:

	PPI data			*NIPS data*		
Model	*Degree*	*ESP*	*Geodesic*	*Degree*	*ESP*	*Geodesic*
EPM	$0.49 \pm .03$	$0.31 \pm .07$	$0.65 \pm .04$	$0.57 \pm .06$	$0.43 \pm .15$	$0.72 \pm .05$
IRM	$0.30 \pm .04$	$0.15 \pm .07$	$0.25 \pm .07$	$0.29 \pm .08$	$0.46 \pm .10$	$0.36 \pm .13$
ER	$0.57 \pm .02$	$0.45 \pm .03$	$0.23 \pm .02$	$\mathbf{0.26 \pm .06}$	$0.69 \pm .06$	$0.41 \pm .06$
ACL	$0.28 \pm .02$	$\mathbf{0.09 \pm .02}$	$0.34 \pm .03$	$0.42 \pm .05$	$0.51 \pm .06$	$0.50 \pm .04$
RW-U	$\mathbf{0.23 \pm .03}$	$0.17 \pm .02$	$\mathbf{0.16 \pm .08}$	$\mathbf{0.26 \pm .04}$	$\mathbf{0.33 \pm .05}$	$\mathbf{0.22 \pm .08}$
RW-SB	$0.27 \pm .02$	$0.11 \pm .02$	$0.34 \pm .04$	$0.42 \pm .05$	$0.39 \pm .06$	$0.45 \pm .07$

For the ER model, the IRM, and the two random walk models, PPD estimates are shown in more detail in Figure 7, on a logarithmic scale. Some comments:

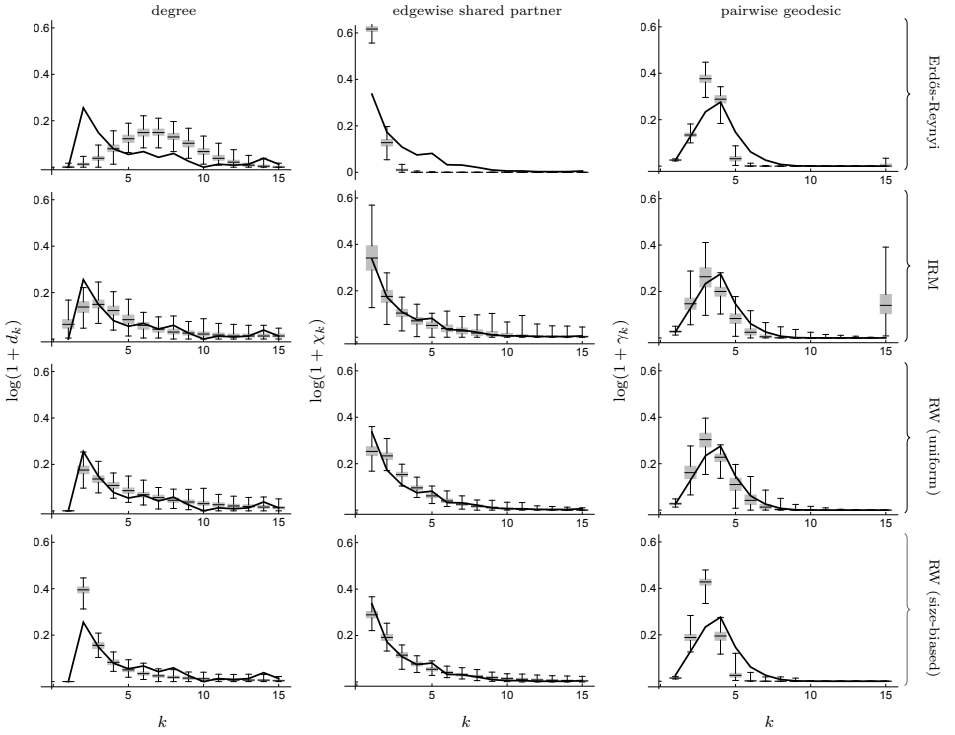

FIG 7. *Estimated PPDs for the PPI data set of three statistics: Normalized degree d_k, normalized edgewise shared partner statistic χ_k, and normalized pairwise geodesic γ_k. Results are shown for four models: Erdős-Rényi (top row), IRM (second row), $\mathbf{RW}_{\mathrm{U}}$ (third row), and $\mathbf{RW}_{\mathrm{SB}}$ (bottom row). The black line represents the distribution from the PPI data.*

- In terms of the protocol of [26], the uniform random walk model provides the best fit on both data sets.
- Of the models used for comparison, the IRM does similarly well, although it does place significant posterior mass on d_0 and γ_∞ (since it tends to generate isolated vertices).
- In the case of the IRM, good fit comes at the price of model complexity: The IRM is a prior on stochastic blockmodels with an infinite number of classes, a finite number of which are invoked to explain a graph of finite size. For the PPI data set, for example, the IRM posterior sharply concentrates at 9 classes, which amounts to 53 scalar parameters, compared to 2 parameters of the RW model.

The numerical results can be illustrated by sampling reconstructions of the input network from models fitted to the data. Figure 8 compares reconstructions of the PPI data set generated by the random walk models to those generated by the IRM and the ER model. Such visual network comparisons should be treated with great caution; nonetheless, the comparison underscores that for some types of data,

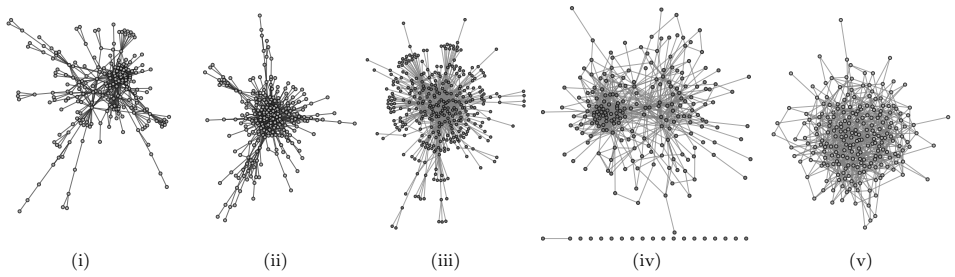

FIG 8. *Reconstructions of the PPI network (i), sampled from the posterior mean of (ii) the* $\mathbf{RW}_{\mathrm{U}}$ *model, (iii) the* $\mathbf{RW}_{\mathrm{SB}}$ *model, (iv) the IRM, and (v) the Erdős-Renyi model.*

including the data set depicted here, the random walk model provides an arguably better structural fit.

5.4. *Latent order and vertex centrality.* Various network statistics attempt to measure the importance of individual vertices for interactions within the network, a property known as **centrality**. This idea can be formalized in a variety of ways. A simple measure of centrality is simply the degree of a vertex. Other examples include Eigenvalue Centrality, Katz Centrality, and Information Centrality [e.g. 37, 29], which each measure different properties of the graph. Under a random walk model, the number of random walks passing through a vertex in the process of network formation provides an obvious measure of that vertex's importance to the formation of the network. A simpler proxy is the order in which vertices are inserted. If only the final graph G_T is observed, this is the latent vertex order $S_{1:T}$ in Theorem 3.3, which is in turn determined by the latent edge order Π, as given by the posterior distribution of the bridge $G_{1:T}$ generated by the inference algorithm. The median of the order of vertices in Π under the $\mathbf{RW}_{\mathrm{U}}$ model is listed below for the NIPS dataset, for those nine authors with earliest median appearance. The sampling order is listed as the vertex's rank, with the numerical value in parentheses. Vertex ranks in other measures of centrality are also listed, with numerical values in parentheses.

Sampling Order (median order)			*Degree (degree)*	*Btwn. Cent. (# shortest paths)*	*Info. Cent. (info. cent.)*	*Katz Cent. (Katz cent.)*
●	(r)	1 (6)	14 (4)	2 (2376)	1 (0.4478)	6 (0.2380)
●	(l)	1 (6)	4 (8)	3 (2340)	2 (0.4443)	2 (0.3740)
●	(r)	2 (7)	1 (12)	1 (2797)	3 (0.4374)	3 (0.2571)
●	(l)	2 (7)	8 (6)	5 (2116)	5 (0.4178)	4 (0.2539)
●		3 (8)	2 (10)	11 (728)	4 (0.4259)	1 (0.3789)
●		4 (10)	2 (10)	4 (2196)	11 (0.3877)	5 (0.2458)
●		5 (12)	14 (4)	6 (1870)	17 (0.3645)	35 (0.0667)
●		6 (13)	14 (4)	16 (173)	6 (0.4040)	19 (0.1291)
●		7 (14)	14 (4)	23 (28)	7 (0.3956)	7 (0.2264)

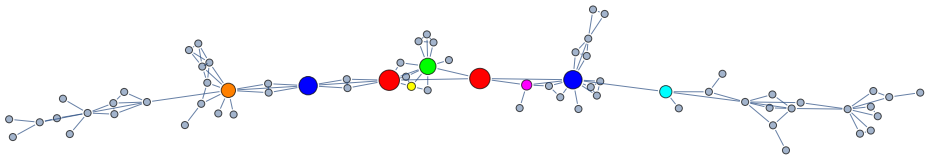

FIG 9. *NIPS authors sampled earliest in latent sequence under the* $\mathbf{RW}_{\mathrm{U}}$ *model.*

Figure 9 maps these vertices in the graph. Since one would generally expect that vertices inserted into the graph earlier tend to have larger degree, it is interesting to note that S seems to correlate more closely with Betweenness Centrality and Information Centrality than with vertex degree.

6. Discussion and open questions. We have introduced a class of network models that can account for dependence of new links on existing links, but are nonetheless tractable. In light of our results, there are two reasons for tractability:

- The effects of the parameters α and λ are sufficiently distinct that inference is possible—a single graph generated by the model carries sufficient signal for parameter values to be recovered with high accuracy, as evident in Figure 5.
- There exists a latent variable—the edge insertion order Π—conditioning on which greatly simplifies the distribution of $G_{1:T} \,|\, G_T$.

Conditioning on history information, either on the order in which vertices are inserted, or on entire history Π, is instrumental both for theoretical results and for inference. As mentioned in the introduction, a qualitatively similar effect—the joint distribution of the network graph simplifies conditionally on a suitable latent variable—is observed in other network models, such as graphon models and the configuration model.

The inference algorithms in Section 4 work well for moderately sized graphs (of up to a few hundred vertices). The dominant source of computational complexity is sampling orders of vertices—similar to inference algorithms for graphon models (which have comparable input size limitations), although the latent order plays a very different role here than in an exchangeable model. Although improvements in efficiency or approximations are certainly possible, we do not expect the methods to scale to very large graphs. Such graphs are not the subject of this work. It is crucial that, under the random walk model, graphs of feasible size do provide sufficient data for reliable inference. For more complicated models, that might not be the case.

To conclude, we note there are open questions beyond the scope of this paper. Some of these are obvious, such as applications to dynamic network problems; generalizations to disconnected graphs (e.g. using a disconnected seed graph); or which choices of the distribution P of the random walk length in Theorem 3.2 satisfy condition (12), and, more generally, for which choices of P the model is tractable, aside from Poisson and negative binomial distributions.

A question we consider intriguing is whether the random walk itself results in power-law behavior. Intuitively, the probability to reach a vertex by random walk increases with degree; that suggests the random walk might result in an effect similar to preferential attachment. For any $\mathbf{RW}(\alpha, \text{Poisson}_+(\lambda))$ model, the probability (1) can be written as

$$\mathbb{P}(V_{end} = v | V_0 = u,\, G) = \left[\tfrac{d_v}{\text{vol}(G)} + \sqrt{\tfrac{d_v}{d_u}} \sum_{i=2}^{n} (1-\sigma_i) e^{-\lambda\sigma_i} \boldsymbol{\psi}_i(u) \boldsymbol{\psi}_i'(v)\right],$$

where $(\boldsymbol{\psi}_i)$ are the eigenvectors, and (σ_i) the eigenvalues, of the normalized graph Laplacian $\mathbf{L}$. Thus, the probability to terminate at v is a mixture of a preferential attachment term (proportional to d_v), and a term depending on λ and on the structure of G at various scales. Our power law results in Section 3 stem from size-biased selection from the vertex set, not the random walk. This mechanism is not used in the $\mathbf{RW}_{\text{U}}$ model on multigraphs, which empirically nonetheless exhibits a heavy-tailed degree distribution. That seems to suggest an affirmative answer, but at present, we have no proof.

APPENDIX A: PROOFS

A.1. Proof of Proposition 3.1. If K has law $\text{Poisson}_+(\lambda)$, (1) becomes

$$\begin{aligned} Q^\lambda &= \sum_{k=0}^{\infty} \frac{e^{-\lambda}\lambda^k}{k!} (D^{-1}A)^{k+1} = D^{-1/2} \left[(\mathbb{I}_n - \mathbf{L})\, e^{-\lambda} \sum_{k=0}^{\infty} \frac{\lambda^k\, (\mathbb{I}_n - \mathbf{L})^k}{k!} \right] D^{1/2} \\ &= D^{-1/2} \left[(\mathbb{I}_n - \mathbf{L})\, e^{-\lambda\mathbf{L}} \right] D^{1/2} = D^{-1/2} \left[(\mathbb{I}_n - \mathbf{L})\, \mathbf{K}^\lambda \right] D^{1/2}. \end{aligned}$$

Since the spectral norm of $\mathbb{I} - \mathbf{L}$ is 1, series is absolutely convergent. For $K \sim \text{NB}_+(r,p)$,

$$\sum_{k=0}^{\infty} \frac{\Gamma(k+r)}{\Gamma(k+1)\Gamma(r)} p^k (1-p)^r (D^{-1}A)^{k+1} = D^{-1/2} \left[(\mathbb{I}_n - \mathbf{L}) \left(\mathbb{I}_n + \frac{p}{1-p}\mathbf{L} \right)^{-r} \right] D^{1/2}$$

similarly yields $Q^{r,p} = D^{-1/2} \left[(\mathbb{I}_n - \mathbf{L})\, \mathbf{K}^{r,p} \right] D^{1/2}$.

A.2. Proof of Theorem 3.2. We begin with a lemma used in the proofs of Theorem 3.2 and Theorem 3.3. Let $\delta_j^{(1)}(t)$ be an indicator variable that is equal to 1 if the first end of the t-th edge is attached to vertex v_j, and 0 otherwise; likewise define $\delta_j^{(2)}(t)$ for the second end of the t-th edge.

LEMMA 1. *Suppose a sequence of graphs is generated with law $\mathbf{RW}_{\text{SB}}(\alpha, P)$, such that P is a probability distribution on $\mathbb{N}$. Then the size-biased distribution $\mathbb{S}_t$ is left-invariant under the mixed random walk probability Q_t induced by P, that is*

$$\mathbb{S}_t' Q_t = \mathbb{S}_t' \qquad \text{where } [Q_t]_{uv} := \mathbb{P}\{V_{end} = v \,|\, V_0 = u, G_t\} \tag{29}$$

holds for every $t \in \mathbb{N}_+$. Furthermore, for each $v_j \in \mathcal{V}(G_t)$, the following holds:

$$\mathbb{E}[\delta_j^{(1)}(t+1) \,|\, G_t] = \frac{1}{1-\alpha} \mathbb{E}[\delta_j^{(2)}(t+1) \,|\, G_t] = \frac{\deg_t(v_j)}{2t}. \tag{30}$$

PROOF. Substituting (1) into $\mathbb{S}_t' Q_t$ yields

$$\mathbb{S}_t' Q_t = \sum_k \mathbb{S}_t' (D^{-1}A)^k P\{K = k\} = \sum_k \mathbb{S}_t' P\{K = k\} = \mathbb{S}_t' \,. \tag{31}$$

The $\mathbf{RW}_{\text{SB}}$ model samples the first end of each edge from the size-biased distribution:

$$\mathbb{E}[\delta_j^{(1)}(t+1) \,|\, G_t] = \mathbb{P}[V_{t+1} = v_j \,|\, G_t] = \frac{\deg_t(v_j)}{2t} \,. \tag{32}$$

For the second end, denote by $\mathbb{1}_{j,t}$ the indicator vector for vertex v_j. Then

$$\mathbb{E}[\delta_j^{(2)}(t+1) \,|\, G_t] = \sum_u \mathbb{P}[V_{t+1} = u \,|\, G_t] \big[Q_t\big]_{uv} = (1-\alpha)\mathbb{S}_t' Q_t \mathbb{1}_{j,t} = (1-\alpha)\frac{\deg_t(v_j)}{2t} \,.$$

□

In the setup of Theorem 3.2, asymptotic expected degree counts are as follows:

LEMMA 2. *Let a sequence of multigraphs $(G_1, G_2, \dots)$ have law $\mathbf{RW}_{\text{SB}}(\alpha, P)$, for a probability distribution P on $\mathbb{N}$ satisfying Equation (12). Then*

$$\frac{\mathbb{E}[m_{d,t}]}{\alpha t} \xrightarrow{a.s.} p_d \qquad \text{for each } d \in \mathbb{N} \text{ as } t \to \infty \,, \tag{33}$$

where p_d is given in (10).

PROOF. Equation (12) yields the recurrence

$$\mathbb{E}(m_{d,t+1}) = \alpha \cdot \delta_1(d) + \mathbb{E}(m_{d,t})\Big(1 - \frac{(2-\alpha)d}{2t}\Big) + \mathbb{E}(m_{d-1,t})\frac{(2-\alpha)(d-1)}{2t} + o(1) \,,$$

with $m_{0,t} = 0$ for all t. This can be written more generally as

$$M(d, t+1) = (1 - b(t)/t)M(d,t) + g(t) \,. \tag{34}$$

If $b(t) \to b$ and $g(t) \to g$, then $M(d,t)/t \to g/(b+1)$, by [18, Lemmas 4.1.1, 4.1.2]. For $d = 1$, we have $b(t) = b = (2-\alpha)/2$ and $g(t) = g = \alpha$, so $\mathbb{E}(m_{1,t})/t \to \frac{2\alpha}{4-\alpha} = \alpha p_1$. Proceeding by induction, $b = (2-\alpha)d/2$ and $g = \alpha p_{d-1}(2-\alpha)(d-1)/2$ yield

$$\frac{\mathbb{E}[m_{d,t}]}{\alpha t} \quad \to \quad p_{d-1}\tfrac{(2-\alpha)(d-1)}{2+(2-\alpha)d} = \frac{2}{4-\alpha}\prod_{j=1}^{d-1}\frac{j}{\frac{2}{2-\alpha}+j+1} = \frac{2}{2-\alpha}\frac{\Gamma(d)\Gamma(2+\frac{\alpha}{2-\alpha})}{\Gamma(d+2+\frac{\alpha}{2-\alpha})} \,,$$

for $d > 1$, which is just p_d as defined in (10). □

The following result, the proof of which can be found in [14, Section 3.6], shows that the random variable $m_{d,t}$ concentrates about its mean:

LEMMA 3. *Let a sequence of multigraphs* $(G_1, G_2, \dots)$ *have law* $\mathbf{RW}_{\text{SB}}(\alpha, P)$, *for a probability distribution* P *on* $\mathbb{N}$ *satisfying Equation* (12). *Then*

$$\mathbb{P}\Big[\ \Big|\frac{m_{d,t}}{\alpha t} - p_d\ \Big| \leq 2\sqrt{\frac{d^3 \log t}{\alpha^2 t}}\ \Big] \geq 1 - 2(t+1)^{d-1}t^{-d} = 1 - o(1)\ . \tag{35}$$

PROOF OF THEOREM 3.2. Lemma 3 shows that $\frac{m_{d,t}}{\alpha t}$ concentrates around its mean, and with Lemma 2 showing that the mean is p_d, the proof of the theorem is complete. □

A.3. Proof of Theorem 3.3. By definition of **RW** models (see Algorithm 1), at each step t a new vertex appears with probability α, which is represented as a collection of Bernoulli(α) random variables $(B_t)_{t\geq 2}$. Alternatively, there is the collection of steps $S_1, S_2, \dots, S_j, \dots$ at which new vertices appear, such that S_j is the step at which the j-th vertex appears. By the fundamental relationship between Bernoulli and Geometric random variables, the sequence $S_1, S_2, \dots, S_j, \dots$ can be sampled independently of the graph sequence as $S_1 = S_2 = 1$ and

$$S_j = S_{j-1} + \Delta_j, \text{ where } \Delta_j \overset{\text{iid}}{\sim} \text{Geometric}(\alpha), \text{ for } j > 2. \tag{36}$$

In what follows, we condition on the sequence $(S_j)_{j\geq 1}$ unless explicitly stated otherwise. We begin by calculating the expected degree of the j-th vertex.

LEMMA 4. *Let a sequence of multigraphs* $(G_1, G_2, \dots)$ *have law* $\mathbf{RW}_{\text{SB}}(\alpha, P)$ *such that Lemma 1 applies. Let* $d_j(t) := \deg_t(v_j)$ *be the degree of* v_j *in graph* G_t, *where* v_j *is the* j*-th vertex to appear in the graph sequence, and let* $\rho = 1 + \frac{\alpha}{2-\alpha}$. *Then conditional on* S_j, $d_j(t)t^{-1/\rho}$ *converges almost surely to a random variable* ξ_j *as* $t \to \infty$, *and*

$$\mathbb{E}[d_j(t) \,|\, S_j] = \frac{\Gamma(S_j)\Gamma(t + \frac{1}{\rho})}{\Gamma(S_j + \frac{1}{\rho})\Gamma(t)} \tag{37}$$

PROOF. Let $\mathcal{A}_t$ denote the σ-algebra generated after t steps. Then for $t \geq S_j$,

$$\mathbb{E}[d_j(t+1) \,|\, \mathcal{A}_t] = d_j(t) + \mathbb{E}[\delta_j^{(1)}(t+1) \,|\, \mathcal{A}_t] + \mathbb{E}[\delta_j^{(2)}(t+1) \,|\, \mathcal{A}_t] = d_j(t)(1 + \tfrac{1}{\rho t})\ ,$$

where the final identity follows from Lemma 1, and $d_j(S_j) = 1$. The sequence

$$M_j(t) := d_j(t)\frac{\Gamma(S_j + \frac{1}{\rho})}{\Gamma(S_j)} \cdot \frac{\Gamma(t)}{\Gamma(t + \frac{1}{\rho})} \qquad \text{for } t \geq S_j$$

is a non-negative martingale with mean 1, by (A.3). Therefore, $M_j(t)$ converges almost surely to a random variable M_j as $t \to \infty$. Taking expectations on both sides and rearranging (A.3) yields (37). With Stirling's formula,

$$\frac{d_j(t)}{t^{1/\rho}} \xrightarrow{\text{a.s.}} M_j \frac{\Gamma(S_j)}{\Gamma(S_j + \frac{1}{\rho})} := \xi_j\ . \tag{38}$$

□

We now consider the joint distribution of the limiting random variables $(\xi_j)_{j\geq 1}$. The approach, which is adapted from [35] (see also [18, 43]), is to analyze a martingale that yields the moments of $(\xi_j)_{j\geq 1}$. First, define for $t \geq S_j$,

$$R_{j,k}(t) := \frac{\Gamma(d_j(t)+k)}{\Gamma(d_j(t))\Gamma(k+1)} , \tag{39}$$

At a high level, $R_{j,k}(t) \approx d_j(t)^k/k!$ for large t, so (38) shows that properly scaled $R_{j,k}(t)$ should converge to $\xi_j^k/k!$ for each j. We make this precise in what follows.

Let $\vec{j} := (j_i)$ be an ordered collection of vertices such that $1 \leq j_1 < j_2 < \cdots < j_r$, and let $\vec{k} := (k_i)$ be a corresponding vector of moments. To define a suitable martingale, we abbreviate

$$\mu_{\vec{k}} := \sum_{i=1}^{r} k_i \quad \text{and} \quad \Sigma_{\vec{j},\vec{k}}(t) := \sum_{i,l=1}^{r} \left(\frac{1}{2}\right)^{\mathbb{1}_{i=l}} \frac{k_i(k_l - \mathbb{1}_{i=l})}{d_{j_l}(t) + \mathbb{1}_{i=l}} [Q_t]_{j_i j_l} ,$$

and note $[Q_t]_{jj'} = 0$ for all $t < \max\{S_j, S_{j'}\}$. Define

$$c_{\vec{j},\vec{k}}(t) = \Gamma(t)\Big[\prod_{s=0}^{t-1}\big(s + \tfrac{\mu_{\vec{k}}}{\rho} + \tfrac{1-\alpha}{2}\Sigma_{\vec{j},\vec{k}}(s)\big)\Big]^{-1} , \tag{40}$$

which is a random variable since Q_t is random, and is $\mathcal{A}_t$-measurable for each t. Since $[Q_t]_{j_i j_l} < 1$, and since $d_j(t) \xrightarrow{\text{a.s.}} \infty$ as $t \to \infty$ by by Lemma 4, $\Sigma_{\vec{j},\vec{k}}(t) \xrightarrow{\text{a.s.}} 0$ as $t \to \infty$, which yields

$$c_{\vec{j},\vec{k}}(t) = \frac{\Gamma(t)}{\Gamma(t + \frac{1}{\rho}\mu_{\vec{k}})}(1 + o(1)) = t^{-\mu_{\vec{k}}/\rho}(1 + o(1)) \quad \text{as } t \to \infty . \tag{41}$$

Note that

$$\lim_{t\to\infty} c_{\vec{j},\vec{k}}(t) \prod_{i=1}^{r} R_{j_i,k_i} = \prod_{i=1}^{r} \frac{\xi_{j_i}^{k_i}}{\Gamma(k_i+1)} , \tag{42}$$

if the limit exists. Existence is based on the following result:

LEMMA 5. *Let $\mathcal{A}_t$ denote the σ-algebra generated by a $\mathbf{RW}_{\text{SB}}(\alpha, P)$ multigraph sequence up to step t. Let $r > 0$ and $1 \leq j_1 < j_2 < \cdots < j_r$ be integers, and real-valued $k_1, k_2, \ldots, k_r > -1$. Then with $R_{j,k}(t)$ defined in (39) and $c_{\vec{j},\vec{k}}(t)$ defined in (40),*

$$Z_{\vec{j},\vec{k}}(t) := c_{\vec{j},\vec{k}}(t) \prod_{i=1}^{r} R_{j_i,k_i}(t) \tag{43}$$

is a nonnegative martingale for $t \geq \max\{S_{j_r}, 1\}$. If $k_1, k_2, \ldots, k_r > -\frac{1}{2}$, then $Z_{\vec{j},\vec{k}}(t)$ converges in L_2.

PROOF. It can be shown that

$$\mathbb{E}\Big[\prod_{i=1}^{r} R_{j_i,k_i}(t+1) \;\Big|\; \mathcal{A}_t\Big] = \Big(\prod_{i=1}^{r} R_{j_i,k_i}(t)\Big)\Big(1 + \frac{\mu_{\vec{k}}}{\rho t} + \frac{1-\alpha}{2t}\Sigma_{\vec{j},\vec{k}}(t)\Big)\,, \tag{44}$$

from which it follows that

$$\mathbb{E}[Z_{\vec{j},\vec{k}}(t+1)\,|\,\mathcal{A}_t] = c_{\vec{j},\vec{k}}(t+1)R_{\vec{j},\vec{k}}(t)\Big(1+\frac{\mu_{\vec{k}}}{\rho t}+\frac{1-\alpha}{2t}\Sigma_{\vec{j},\vec{k}}(t)\Big) = c_{\vec{j},\vec{k}}(t)R_{\vec{j},\vec{k}}(t) = Z_{\vec{j},\vec{k}}(t)\,. \tag{45}$$

Furthermore, $Z_{\vec{j},\vec{k}}(\max\{S_{j_r},1\}) > 0$. By (41) and properties of the gamma function,

$$Z_{\vec{j},\vec{k}}(t)^2 \leq Z_{\vec{j},2\vec{k}}(t)\prod_{i=1}^{r}\binom{2k_i}{k_i}\,. \tag{46}$$

By (45), $Z_{\vec{j},2\vec{k}}(t)$ is a martingale with finite expectation for $2k_1,\ldots,2k_r > -1$. Therefore, $Z_{\vec{j},\vec{k}}(t)$ is an L_2-bounded martingale and hence converges in L_2 (and also in L_1) for $k_1,\ldots,k_r > -\frac{1}{2}$. □

Combining the auxiliary results above, we can now give proof of the result.

PROOF OF THEOREM 3.3. The limit of $Z_{\vec{j},\vec{k}}(t)$ is (42), which enables calculation of the moments of ξ_j. In particular, for any vertex v_j, $j \in \mathbb{N}_+$, $k \in \mathbb{R}$ such that $k > -\frac{1}{2}$,

$$\mathbb{E}\Big[\tfrac{\xi_j^k}{\Gamma(k+1)} \;\Big|\; S_j\Big] = \lim_{t\to\infty}\mathbb{E}[Z_{j,k}(t)\,|\,S_j] = \mathbb{E}[Z_{j,k}(S_j)\,|\,S_j] = \tfrac{\Gamma(S_j)}{\Gamma(S_j+\frac{k}{\rho})}\,. \tag{47}$$

Although the joint moments involve $\Sigma_{\vec{j},\vec{k}}$, the moments (47) characterize the marginal distribution of $\xi_j\,|\,S_j$ via the Laplace transform. Since

$$\mathbb{E}[\xi_j^k\,|\,S_j] = \frac{\Gamma(S_j)\Gamma(k+1)}{\Gamma\big(S_j+\frac{k}{\rho}\big)} = \frac{\Gamma\left(\rho(S_j-1)+1+k\right)\Gamma(S_j)}{\Gamma\left(\rho(S_j-1)+1\right)\Gamma\big(S_j+\frac{k}{\rho}\big)}\frac{\Gamma\left(\rho(S_j-1)+1\right)\Gamma(k+1)}{\Gamma\left(\rho(S_j-1)+1+k\right)}\,,$$

$\mathbb{E}[\xi_j^k\,|\,S_j] = \mathbb{E}[M_j^k B_j^k\,|\,S_j]$, where M_j is a generalized Mittag-Leffler(ρ^{-1}, S_j-1) variable [27], and B_j is $B_j \sim \mathrm{Beta}(1,\rho(S_j-1))$. It follows that $M_j \perp\!\!\!\perp_{S_j} B_j$.

It remains to show (16). We begin by noting that the left-hand side of (47) is an expectation that conditions on S_j. Define $\tilde{S}_j := S_j - 1 - (j-2)$, which is marginally distributed as $\mathrm{NB}(j-2, 1-\alpha)$. Then

$$\begin{aligned}
\mathbb{E}\big[\mathbb{E}[\xi_j^k\,|\,S_j]\big] &= \mathbb{E}\Big[\frac{\Gamma(\tilde{S}_j+1+(j-2))\Gamma(k+1)}{\Gamma(\tilde{S}_j+1+(j-2)+\frac{k}{\rho})}\Big] \\
&= \sum_{t=0}^{\infty}\frac{\Gamma(t+1+(j-2))\Gamma(k+1)}{\Gamma(t+1+(j-2)+\frac{k}{\rho})}\frac{\Gamma(t+j-2)}{\Gamma(j-2)\Gamma(t+1)}(1-\alpha)^t\alpha^{j-2} \qquad (48)\\
&= \frac{\Gamma(k+1)\Gamma(j-1)}{\Gamma(j-1+\frac{k}{\rho})}\alpha^{\frac{k}{\rho}}\,{}_2F_1(1+\frac{k}{\rho},\frac{k}{\rho};j-1+\frac{k}{\rho};1-\alpha)\,,
\end{aligned}$$

where ${}_2F_1(a,b;c;z)$ is the ordinary hypergeometric function. For $j \to \infty$,

$$\lim_{j\to\infty} \mathbb{E}\big[\mathbb{E}[\xi_j^k \,|\, S_j]\big] = \lim_{j\to\infty} \tfrac{\Gamma(k+1)\Gamma(j-1)}{\Gamma(j-1+\frac{k}{\rho})}\alpha^{\frac{k}{\rho}}(1+O(j^{-1})) = \Gamma(k+1)\alpha^{\frac{k}{\rho}} j^{-\frac{k}{\rho}}(1+O(j^{-1}))\ .$$

follows using the series expansion of ${}_2F_1(a,b;c;z)$. □

A.4. Proof of Proposition 3.4.

Proof of Proposition 3.4. Existence of the limit as $\lambda \to \infty$ follows from the existence of the limiting (stationary) distribution for each $t \in \mathbb{N}_+$. For equivalence, it suffices to show that given any connected graph G, the conditional distribution over graphs G' is the same for $\mathbf{RW}_{\text{SB}}(\alpha,\infty)$ and $\mathbf{ACL}(\alpha)$. The probability of attaching a new vertex to an existing vertex v in both models is $\alpha\deg(v)/\mathrm{vol}(G)$. With probability $1-\alpha$, a new edge is inserted between existing vertices u and v, and so it remains to show that in this case the distribution over pairs of vertices, $\nu(u,v)$, is the same. We have

$$\nu(u,v) = 2\frac{\deg(u)}{\mathrm{vol}(G)}\frac{\deg(v)}{\mathrm{vol}(G)} \quad \text{and} \quad \nu(u,v) = \frac{\deg(u)}{\mathrm{vol}(G)}[Q^\lambda]_{uv} + \frac{\deg(v)}{\mathrm{vol}(G)}[Q^\lambda]_{vu}$$

for the ACL and $\mathbf{RW}_{\text{SB}}$ model respectively, with Q^λ. First, consider $(\mathbb{I}_n - \mathbf{L})\mathbf{K}^\lambda$ in terms of the spectrum of $\mathbf{L}$. Let $0 = \sigma_1 \le \cdots \le \sigma_n \le 2$ be the eigenvalues for a graph on n vertices, with eigenvectors $\boldsymbol{\psi}_i$. Then

$$(\mathbb{I}_n - \mathbf{L})\mathbf{K}^\lambda = (\mathbb{I}_n - \mathbf{L})e^{-\lambda\mathbf{L}} = \sum_{i=1}^n (1-\sigma_i)e^{-\lambda\sigma_i}\boldsymbol{\psi}_i\boldsymbol{\psi}_i'\ . \tag{49}$$

When $\lambda \to \infty$, only the eigenvector $\boldsymbol{\psi}_1$ corresponding to the eigenvalue $\sigma_1 = 0$ contributes to the random walk probabilities, i.e. $\lim_{\lambda\to\infty}(\mathbb{I}_n - \mathbf{L})e^{-\lambda\mathbf{L}} = \boldsymbol{\psi}_1\boldsymbol{\psi}_1'$. The limit satisfies $\boldsymbol{\psi}_1 \propto D^{1/2}\mathbb{1}$, where $\mathbb{1}$ is the vector of all ones [13, Ch. 2]. Therefore,

$$\lim_{\lambda\to\infty}[Q^\lambda]_{uv} = \big[\mathbb{1}\mathbb{1}'D\tfrac{1}{\mathrm{vol}(G)}\big]_{uv} = \frac{\deg(v)}{\mathrm{vol}(G)}\ ,$$

and the result follows. □

A.5. Proofs for Section 4.

Proof of Proposition 4.1. Proposition 2.2 of Del Moral and Murray [16] shows that

$$\mathbb{E}[f(G_{1:T})\,|\,G_1,G_T] = L_\theta^1(G_1)^{-1}\mathbb{E}\big[f(G_{1:T})L_\theta^{T-1}(G_{T-1})\tfrac{h_t(G_t)}{h_t(G_{T-1})}\textstyle\prod_{s=2}^{T-1}\frac{h_s(G_s)}{h_{s-1}(G_{s-1})}\,|\,G_1\big]\ ,$$

for any approximation function h_t that is positive outside a null set. For our choice of h_t, the proposal density (22) only places probability mass on graphs with non-zero bridge likelihood, making $\mathbb{1}\{h_t = 0\}$ a probability zero event. For an SMC

approximation to be consistent, the target density must be absolutely continuous with respect to the proposal density at each step t [see, e.g. 40, 17]. In other words, the proposal density must be a valid importance sampling distribution at each step. The target, $\mathcal{L}_\theta(G_{1:t} \,|\, G_T, G_1)$ is absolutely continuous with respect to $\prod_{s=2}^{t} r_\theta(G_s \,|\, G_{s-1})$ from (22) by construction if properties (P1) and (P2) hold, and the result follows. □

Proof of Proposition 4.2. Theorem 4 in [3], of which Proposition 4.2 is a special case, requires that: (a) The resampling scheme in the SMC algorithm is unbiased, i.e. each particle is resampled with probability proportional to its weight. (b) $\mathcal{L}_\theta(G_{1:t} \,|\, G_T, G_1)$ is absolutely continuous with respect to $\prod_{s=2}^{t} r_\theta(G_s \,|\, G_{s-1})$ for any θ. Condition (a) is satisfied by the multinomial, residual, and stratified resampling schemes. As discussed in the proof of Proposition 4.1, the condition (b) holds by construction if properties (P1) and (P2) hold. □

APPENDIX B: THE PARTICLE GIBBS UPDATES

The particle Gibbs sampler updates are performed in three blocks: (α, λ), $(\mathbf{B}, \mathbf{K})$, and $G_{2:(T-1)}$. As we describe in the following subsections, we make use of conditional conjugacy and tools from spectral graph theory to increase sampling efficiency. We discuss each block of updates in turn.

Updating the parameters α and λ. The model is specified in terms of distributions for the latent variables, and placing conjugate priors on their parameters we have,

$$B_t \,|\, \alpha \overset{\text{iid}}{\sim} \text{Bernoulli}(\alpha), \quad \alpha \sim \text{Beta}(a_\alpha, b_\alpha) \tag{50}$$

$$K_t \,|\, \lambda \overset{\text{iid}}{\sim} \text{Poisson}_+(\lambda), \quad \lambda \sim \text{Gamma}(a_\lambda, b_\lambda). \tag{51}$$

The Gibbs updates are hence $\alpha \,|\, \mathbf{B} \sim \text{Beta}\,(\chi, \omega)$ and $\lambda \,|\, \mathbf{K} \sim \text{Gamma}\,(\kappa, \tau)$, with

$$\chi := a_\alpha + \sum_{t=2}^{T} B_t, \quad \omega := b_\alpha + (T-1) - \sum_{t=2}^{T} B_t \tag{52}$$

$$\kappa := a_\lambda + \sum_{t=2}^{T} K_t, \quad \tau := b_\lambda + (T-1)\,. \tag{53}$$

Updating the latent variables B and K. The sequential nature of the model allows each of the latent variables B_t and K_t to be updated individually. We use conditional independence to marginalize out sampling steps where possible, since such marginalization increase the sampling efficiency [44]. For both $\mathbf{B}$ and $\mathbf{K}$, we marginalize twice: The first transforms the conditional distributions $P(B_t \,|\, \alpha)$ and $P(K_t \,|\, \lambda)$ into the predictive distributions $P(B_t \,|\, \mathbf{B}_{-t})$, and $P(K_t \,|\, \mathbf{K}_{-t})$, which yields

$$B_t \,|\, \mathbf{B}_{-t} \sim \text{Bernoulli}\Big(\frac{\chi_{-t}}{\chi_{-t} + \omega_{-t}}\Big) \quad \text{and} \quad K_t \,|\, \mathbf{K}_{-t} \sim \text{NB}_+\Big(\kappa_{-t}, \frac{1}{1+\tau_{-t}}\Big)\,.$$

The second improvement is made by marginalizing the conditional dependence of B_t on K_t, and of K_t on B_t. That requires some notation: $\delta_t(V_t, U_t)$ indicates that a new edge is added to the graph between vertices V_t and U_t. When a new vertex is attached to V_t, we write $\delta_t(V_t, u^*)$. We abbreviate $\bar{\tau} := \frac{1}{1+\tau}$, and write $\widetilde{\mathcal{N}}(v)$ for the $\{0,1\}$-ball of vertex v, i.e. v and its neighbors. The updates for $\mathbf{B}$ are

$$P(B_t = 1 \,|\, G_{(t-1):t}, \mathbf{K}_{-t}, \mathbf{B}_{-t}) \propto \frac{\delta_t(V_t, u^*)\mu_{t-1}(V_t)\chi_{-t}}{\chi_{-t} + \omega_{-t}}$$
$$P(B_t = 0 \,|\, G_{(t-1):t}, \mathbf{K}_{-t}, \mathbf{B}_{-t}) \propto \delta(V_t, U_t) + \frac{\delta(V_t, u^*)\omega_{-t}\mu_{t-1}(V_t)}{\chi_{-t} + \omega_{-t}} \sum_{u \in \widetilde{\mathcal{N}}(V_t)} \left[Q_{t-1}^{\kappa_{-t}, \bar{\tau}_{-t}}\right]_{V_t, u}$$

with $Q_{t-1}^{\kappa_{-t}, \bar{\tau}_{-t}}$ as in (A.1) with $r = \kappa_{-t}$, $p = \bar{\tau}_{-t}$. The updates for $\mathbf{K}$ are

$$\begin{aligned} P(K_t = k \,|\, G_{(t-1):t}, \mathbf{K}_{-t}, \mathbf{B}_{-t}) &\propto \frac{\Gamma(\kappa_{-t} + k)}{\Gamma(k+1)\Gamma(\kappa_{-t})}(1 - \bar{\tau}_{-t})^{\kappa_{-t}} \bar{\tau}_{-t}^{k} \ \dots \\ \times \ \Big[\ \delta_t(V_t, u^*)\mu_{t-1}(V_t)\Big(\frac{\chi_{-t}}{\chi_{-t} + \omega_{-t}} + \frac{\omega_{-t}}{\chi_{-t} + \omega_{-t}} \sum_{u \in \widetilde{\mathcal{N}}(V_t)} \left[Q_{t-1}^{k+1}\right]_{V_t, u}\Big) \dots \\ + \ \delta_t(V_t, U_t)\frac{\omega_{-t}}{\chi_{-t} + \omega_{-t}}\Big(\mu_{t-1}(V_t)\left[Q_{t-1}^{k+1}\right]_{V_t, U_t} + \mu_{t-1}(U_t)\left[Q_{t-1}^{k+1}\right]_{U_t, V_t}\Big)\Big] \end{aligned}$$

with $Q_{t-1}^{k+1} = D_{t-1}^{-1/2}(\mathbb{I}_{t-1} - \mathbf{L}_{t-1})^{k+1} D_{t-1}^{1/2}$, the probability of a random walk of length $k+1$ from u to v. For implementation, the distribution for K_t must be truncated at some finite k, which can safely be done at three or four times the diameter of G_T: The total remaining probability mass can be calculated analytically, and the mass is placed on larger k is negligible.

Sampling $G_{2:(T-1)}$. Implementation of Algorithm 2 is straight-forward, with one exception: at each SMC step $G_{t-1} \to G_t$, we collapse the dependence on the particular values of B_t, K_t in that sampling iteration so that edges are proposed from the collapsed transition kernel $q_{\alpha,\lambda}^t(G_t \,|\, G_{t-1}, \mathbf{B}_{-t}, \mathbf{K}_{-t})$. Thus, G_t is composed of G_{t-1} plus a random edge e_t sampled as

$$\begin{aligned} P(e_t = (v,u)) \propto \\ \mathbb{1}\{(v,u) = (V_t, u^*)\}\mu_{t-1}(v)\Big(\tfrac{\chi_{-t}}{\chi_{-t}+\omega_{-t}} + \tfrac{\omega_{-t}}{\chi_{-t}+\omega_{-t}}\Big(\sum_{u \in \widetilde{\mathcal{N}}(v)} \left[Q_{t-1}^{\kappa_{-t}, \bar{\tau}_{-t}}\right]_{v,u}\Big)\Big) + \dots \\ \mathbb{1}\{(v,u) = (V_t, U_t)\}\frac{\omega_{-t}}{\chi_{-t} + \omega_{-t}}\left(\mu_{t-1}(v)\left[Q_{t-1}^{\kappa_{-t}, \bar{\tau}_{-t}}\right]_{v,u} + \mu_{t-1}(u)\left[Q_{t-1}^{\kappa_{-t}, \bar{\tau}_{-t}}\right]_{u,v}\right) . \end{aligned}$$

REFERENCES

[1] Aiello, W., Chung, F., and Lu, L. (2002). Random evolution in massive graphs. In J. Abello, P. M. Pardalos, and M. G. C. Resende, editors, *Handbook of Massive Data Sets*, volume II, pages 97–122.

[2] Ambroise, C. and Matias, C. (2012). New consistent and asymptotically normal parameter estimates for random-graph mixture models. *J. R. Stat. Soc. Ser. B. Stat. Methodol.*, **74**(1), 3–35.

[3] Andrieu, C., Doucet, A., and Holenstein, R. (2010). Particle markov chain monte carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, **72**(3), 269–342.
[4] Barabási, A.-L. and Albert, R. (1999). Emergence of scaling in random networks. *Science*, **186**(5439), 509–512.
[5] Batagelj, V. and Mrvar, A. (2006). Pajek datasets. Available at: http://vlado.fmf.uni-lj.si/pub/networks/data/.
[6] Berger, N., Borgs, C., Chayes, J. T., and Saberi, A. (2014). Asymptotic behavior and distributional limits of preferential attachment graphs. *Ann. Probab.*, **42**(1), 1–40.
[7] Bickel, P. J., Chen, A., and Levina, E. (2011). The method of moments and degree distributions for network models. *Ann. Statist.*, **39**(5), 2280–2301.
[8] Borgs, C., Chayes, J. T., Lovász, L., Sós, V. T., and Vesztergombi, K. (2008). Convergent sequences of dense graphs. I. Subgraph frequencies, metric properties and testing. *Adv. Math.*, **219**(6), 1801–1851.
[9] Borgs, C., Chayes, J. T., Cohn, H., and Zhao, Y. (2014). An l^p theory of sparse graph convergence i: limits, sparse random graph models and power law distributions.
[10] Box, G. E. P. (1980). Sampling and bayes' inference in scientific modelling and robustness. *Journal of the Royal Statistical Society. Series A (General)*, **143**(4), pp. 383–430.
[11] Butland, G., Peregrin-Alvarez, J. M., Li, J., Yang, W., Yang, X., Canadien, V., Starostine, A., Richards, D., Beattie, B., Krogan, N., Davey, M., Parkinson, J., Greenblatt, J., and Emili, A. (2005). Interaction network containing conserved and essential protein complexes in escherichia coli. *Nature*, **433**(7025), 531–537.
[12] Caron, F. and Fox, E. B. (2015). Sparse graphs using exchangeable random measures.
[13] Chung, F. (1997). *Spectral Graph Theory (CBMS Regional Conference Series in Mathematics, No. 92)*. American Mathematical Society.
[14] Chung, F. and Lu, L. (2006). *Complex Graphs and Networks*. (CBMS Regional Conference Series in Mathematics). American Mathematical Society.
[15] Chung, F. and Yau, S.-T. (2000). Discrete Green's functions. *J. Combin. Theory Ser. A*, **91**(1-2), 191–214.
[16] Del Moral, P. and Murray, L. M. (2015). Sequential monte carlo with highly informative observations. *SIAM/ASA Journal on Uncertainty Quantification*, **3**(1), 969–997.
[17] Doucet, A. and Johansen, A. M. (2011). A tutorial on particle filtering and smoothing: fifteen years later. *Oxford Handbook of Nonlinear Filtering*, pages 656–704.
[18] Durrett, R. (2006). *Random Graph Dynamics*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, New York, NY, USA.
[19] Gao, C., Lu, Y., and Zhou, H. H. (2015). Rate-optimal graphon estimation. *Ann. Statist.*, **43**, 2624–2652.
[20] Gelman, A., li Meng, X., and Stern, H. (1996). Posterior predictive assessment of model fitness via realized discrepancies. *Statistica Sinica*, **6**(4), 733–807.
[21] Globerson, A., Chechik, G., Pereira, F., and Tishby, N. (2007). Euclidean Embedding of Co-occurrence Data. *The Journal of Machine Learning Research*, **8**, 2265–2295.
[22] Goldenberg, A., Zheng, A. X., Fienberg, S. E., and Airoldi, E. M. (2010). A survey of statistical network models. *Foundations and Trends in Machine Learning*, **2**(2), 129–233.
[23] Griffiths, R. and Tavare, S. (1994a). Simulating probability distributions in the coalescent. *Theoretical Population Biology*, **46**(2), 131 – 159.
[24] Griffiths, R. C. and Tavare, S. (1994b). Ancestral inference in population genetics. *Statist. Sci.*, **9**(3), 307–319.
[25] Hoff, P. D., Raftery, A. E., and Handcock, M. S. (2002). Latent space approaches to social network analysis. *Journal of the American Statistical Association*, **97**(460), 1090–1098.
[26] Hunter, D. R., Goodreau, S. M., and Handcock, M. S. (2008). Goodness of fit of social network models. *Journal of the American Statistical Association*, **103**(481), 248–258.
[27] James, L. F. (2015). Generalized Mittag Leffler distributions arising as limits in preferential attachment models.
[28] Kemp, C., Tenenbaum, J. B., Griffiths, T. L., Yamada, T., and Ueda, N. (2006). Learning systems of concepts with an infinite relational model. In *Proceedings of the 21st National Conference on Artificial Intelligence - Volume 1*, AAAI'06, pages 381–388. AAAI Press.
[29] Kolaczyk, E. D. (2009). *Statistical Analysis of Network Data*. Springer-Verlag New York.
[30] Kondor, R. I. and Lafferty, J. D. (2002). Diffusion kernels on graphs and other discrete input

spaces. In *Proceedings of the Nineteenth International Conference on Machine Learning*, ICML '02, pages 315–322.

[31] Lafferty, J. and Lebanon, G. (2005). Diffusion kernels on statistical manifolds. *J. Mach. Learn. Res.*, **6**, 129–163.

[32] Liggett, T. M. (2005). *Interacting particle systems*. Springer.

[33] Loomis, C. P., Morales, J. O., Clifford, R. A., and Leonard, O. E. (1953). *Turrialba: Social Systems and the Introduction of Change*. The Free Press, Glencoe, Ill.

[34] Lovász, L. and Szegedy, B. (2006). Limits of dense graph sequences. *Journal of Combinatorial Theory, Series B*, **96**(6), 933–957.

[35] Móri, T. F. (2005). The maximum degree of the barabási-albért random tree. *Comb. Probab. Comput.*, **14**(3), 339–348.

[36] Naesseth, C. A., Lindsten, F., and Schön, T. B. (2014). Sequential monte carlo for graphical models. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Weinberger, editors, *Advances in Neural Information Processing Systems 27*, pages 1862–1870.

[37] Newman, M. (2009). *Networks. An Introduction*. Oxford University Press.

[38] Orbanz, P. and Roy, D. M. (2015). Bayesian models of graphs, arrays and other exchangeable random structures. *IEEE Transactions Pattern Analysis and Machine Intelligence*.

[39] Pitman, J. (2006). *Combinatorial Stochastic Processes*, volume 1875 of *Ecole d'Eté de Probabilités de Saint-Flour*. Springer-Verlag Berlin Heidelberg.

[40] Robert, C. P. and Casella, G. (2004). *Monte Carlo Statistical Methods*. Springer New York.

[41] Smola, A. J. and Kondor, I. (2003). Kernels and regularization on graphs. In B. Schölkopf and M. Warmuth, editors, *Proceedings of the Annual Conference on Computational Learning Theory*, Lecture Notes in Computer Science. Springer.

[42] Thorne, T. and Stumpf, M. P. H. (2012). Graph spectral analysis of protein interaction network evolution. *Journal of The Royal Society Interface*.

[43] van der Hofstad, R. (2016). Random graphs and complex networks. Notes accessed March, 2016, http://www.win.tue.nl/~rhofstad/NotesRGCN.html.

[44] van Dyk, D. A. and Park, T. (2008). Partially collapsed gibbs samplers. *Journal of the American Statistical Association*, **103**(482), 790–796.

[45] Wang, J., Jasra, A., and De Iorio, M. (2014). Computational methods for a class of network models. *Journal of Computational Biology*, **21**(2), 141–161.

[46] Wiuf, C., Brameier, M., Hagberg, O., and Stumpf, M. P. H. (2006). A likelihood approach to analysis of network data. *Proceedings of the National Academy of Sciences*, **103**(20), 7566–7570.

[47] Wolfe, P. J. and Olhede, S. C. (2013). Nonparametric graphon estimation. Preprint. http://arxiv.org/abs/1309.5936.

[48] Xu, Z., Tresp, V., Yu, K., and Kriegel, H. P. (2006). Infinite hidden relational models. In *Proceedings of the 22nd International Conference on Uncertainity in Artificial Intelligence (UAI 2006)*, Cambridge, MA, USA. AUAI Press.

[49] Zhou, M. (2015). Infinite edge partition models for overlapping community detection and link prediction. In G. Lebanon and S. V. N. Vishwanathan, editors, *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, volume 38, pages 1135–1143.

Department of Statistics
1255 Amsterdam Avenue
New York, NY-10027, USA
E-mail: benjamin.reddy@columbia.edu,E-mail: porbanz@stat.columbia.edu